

# Studying and Mitigating Biases in Sign Language Understanding Models

**Katherine Atwell**

Northeastern University  
atwell.ka@northeastern.edu

**Danielle Bragg**

Microsoft Research  
danielle.bragg@microsoft.com

**Malihe Alikhani**

Northeastern University  
m.alikhani@northeastern.edu

## Abstract

Ensuring that the benefits of sign language technologies are distributed equitably among all community members is crucial. Thus, it is important to address potential biases and inequities that may arise from the design or use of these resources. Crowd-sourced sign language datasets, such as the ASL Citizen dataset, are great resources for improving accessibility and preserving linguistic diversity, but they must be used thoughtfully to avoid reinforcing existing biases.

In this work, we utilize the rich information about participant demographics and lexical features present in the ASL Citizen dataset to study and document the biases that may result from models trained on crowd-sourced sign datasets. Further, we apply several bias mitigation techniques during model training, and find that these techniques reduce performance disparities without decreasing accuracy. With the publication of this work, we release the demographic information about the participants in the ASL Citizen dataset to encourage future bias mitigation work in this space.

## 1 Introduction

Within the field of natural language processing, sign languages are under-resourced compared to spoken languages, compounded by the fact that most accessible information (e.g. online resources and social media) is written in a spoken language (Desai et al., 2024). Datasets like the ASL Citizen dataset offer significant potential for improving accessibility and preserving the linguistic richness of sign languages, yet their use requires careful consideration to avoid reinforcing existing biases. In this context, our research aims to explore the factors that might influence the performance of models trained on these datasets, particularly when used for dictionary retrieval tasks.

Because sign languages have comparatively fewer resources than spoken languages, identify-

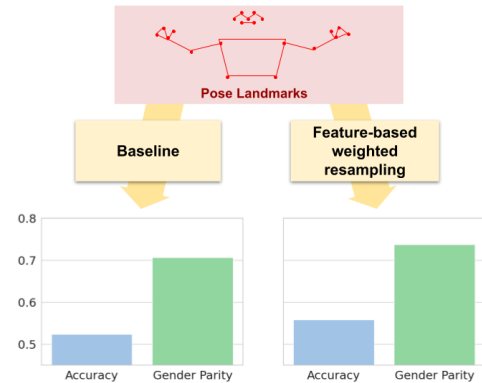


Figure 1: Accuracy and gender parity (calculated by dividing accuracy on female participants by accuracy on male participants) of the baseline pose-based ISLR model released with the ASL Citizen dataset (left) and our best-performing feature-based debiasing technique (right), in which we resample videos with lower video quality scores at a higher rate. Our approach improves both overall model accuracy and the gender parity.

ing biases in existing sign language resources is critical. But biases can manifest differently in sign languages than in spoken languages. For instance, ASL pronouns, unlike English pronouns, are not assigned a gender, so the common method of studying bias in English text through the lens of gendered pronoun use does not apply. Temporal elements, such as signing speed, also come into play, unlike in written language. Signing speed may be impacted by a signer’s fluency, age, etc.

In this work, we analyze how signer demographics and more latent sources of bias may impact models trained on the ASL Citizen dataset for the task of Isolated Sign Language Recognition (ISLR). We first examine the demographic distributions in the ASL Citizen dataset, and present a linguistic analysis of the dataset based on the ASL-Lex (Caselli et al., 2017) annotations for each sign. We then report the prevalence of various linguistic and video-level features among demographics. We examine how demographic features, in conjunction

with lexical and video-level features, may impact model results. Finally, we experiment with multiple debiasing techniques to reduce performance gaps between genders, and find that we are able to reduce these gaps and improve overall model accuracy (Figure 1).

In summary, we present an analysis of demographics, sign-level features, and video-level features in the ASL Citizen dataset and address the following research questions:

1. Which demographic and linguistic factors impact dictionary retrieval results for models trained on the ASL Citizen dataset?
2. Can we use debiasing strategies to mitigate disparate impacts while maintaining high performance for dictionary retrieval models?

With this work, we also release the demographic data for the ASL Citizen dataset<sup>1</sup>, so future researchers can continue to study and mitigate bias in sign language processing systems. Further, we release the code for our experiments and analyses<sup>2</sup>.

## 2 Related Work

Most readily-available information (i.e. online resources and social media) is written, which may limit accessibility for signers. Sign language processing tasks, such as dictionary retrieval, are designed to improve the accessibility of existing systems and resources for Deaf and Hard-of-Hearing (DHH) people. Desai et al. (2024) created the ASL Citizen dataset to supplement existing dictionary retrieval resources with crowd-sourced videos from signers.

The ASL Citizen dataset was released to 1) address the resource gap between sign and spoken languages, and 2) improve *video-based dictionary retrieval* for sign language, where signers demonstrate a particular sign and the system returns a list of similar signs, ranked from most to least similar. Video-based dictionary retrieval systems can help language learners understand the meaning of a sign, and allow signers to access dictionary resources using sign languages (Desai et al., 2024). As a crowd-sourced dataset with videos of individual signs, the ASL Citizen dataset also serves to improve documentation of sign languages. This dataset is the first crowd-sourced dataset of videos

for isolated signs, and members of deaf communities participated in, and were compensated for, this effort. When supplemented with the Sem-Lex benchmark (Kezar et al., 2023a), a crowdsourced ISLR dataset released shortly after, 174k videos in total can be used for ISLR. The ASL Citizen dataset is licensed by Microsoft Research and is bound by the Microsoft Research Licensing Terms<sup>3</sup>.

The ASL Citizen dataset is composed of videos of individual signs for isolated sign language recognition (ISLR). Other ISLR datasets with videos of individual signs have been released, including WL-ASL (Li et al., 2020), Purdue RVL-SLL (Wilbur and Kak, 2006), BOSTON-ASLLVD (Athitsos et al., 2008), and RWTH BOSTON-50 (Zahedi et al., 2005). The above datasets, however, are not crowd-sourced. The closest dataset to the ASL Citizen dataset is the Sem-Lex Benchmark (Kezar et al., 2023a), a crowdsourced ISLR dataset with over 91k videos. Because the Sem-Lex Benchmark does not release demographic information about the participants, we are not able to include it in our bias studies.

The ASL Citizen dataset is made up of crowd-sourced videos from ASL signers, where each video corresponds to a particular sign. The corpus is composed of videos for 2731 unique signs, all of which are contained in the ASL-Lex dataset Caselli et al. (2017), a lexical database of signs with annotations including the relative frequency, iconicity, grammatical class, English translations, and phonological properties of the sign. Thus, researchers studying this dataset can also take advantage of the ASL-Lex annotations. As part of the original data collection effort, demographic information about each participant was collected, but it was not released. With the publication of this work, we release the demographic data in this set, and provide a detailed analysis of this data.

Our analyses of demographics and bias are motivated by evidence in the literature that a signer’s demographics may impact their signing. For instance, characteristics of particular spoken languages or dialects have been shown to influence gestures, and in turn sign production (Cormier et al., 2010). One example of an ASL dialect is Black ASL, which scholarly evidence has shown to be its own dialect (Toliver-Smith and Gentry, 2017), and for which documentation of dialectal differences dates back

<sup>1</sup>Demographics available through the ASL Citizen project page: <https://www.microsoft.com/en-us/research/project/asl-citizen/>

<sup>2</sup><https://github.com/katherine-atwell/mitigating-biases-sign-understanding>

<sup>3</sup>Terms of use at <https://www.microsoft.com/en-us/research/project/asl-citizen/dataset-license/>. We are using this dataset in accordance with its intended use.

to 1965 (Stokoe et al., 1965). Whether an individual speaks Black ASL is likely heavily influenced by their race or ethnicity. An example of geographic differences is Martha’s Vineyard, an island off the coast of the United States, where an entire sign language emerged due to the high prevalence of DHH individuals in this community. Hearing and DHH people alike used this language to communicate until the mid-1900s (Kusters, 2010). There is also a distinct Canadian ASL dialect used by signers in English-speaking areas of Canada (Padden, 2010), which is documented in a dictionary (Bailey et al., 2002). Age of language acquisition also impacts ASL production; delayed first-language acquisition affects syntactic knowledge for ASL signers (Boudreault and Mayberry, 2006) and late acquisition (compared to native acquisition) was found to impact sensitivity to verb agreement (Emmorey et al., 1995).

Previous work also indicates the impact of certain visual and linguistic features on sign language modeling. Training an ISLR model to predict a sign and its phonological characteristics was found to improve model performance by almost 9% (Kezar et al., 2023b). (Sarhan et al., 2023) find improved performance when using attention to focus on hand movements in sign videos.

To our knowledge, there are no existing works that extensively study various sources of model bias on a crowd-sourced dataset of sign videos with labeled participant demographics. With this work, we aim to address this gap with a systematic analysis of the impact of various participant-level, sign-level, and video-level features, and experiment with debiasing techniques to reduce disparities in model performance.

### 3 Data

The ASL Citizen dataset is a crowd-sourced dataset containing 83,399 videos of individual signs in ASL from 52 different participants. The dataset contains 2731 unique signs that are included in the ASL-Lex (Caselli et al., 2017) dataset, a dataset with detailed lexical annotations for each sign. The authors of the original work report some demographic statistics, but the demographics of individual (de-identified) participants have not been released. Here, we provide a detailed report that includes demographic breakdowns and analyses of various linguistic and video features in the dataset, including the breakdown of these features by gen-

der. We release the participant demographics with this work.

#### 3.1 Demographic Distributions

In total, the ASL Citizen dataset is comprised of 32 (61.5%) women and 20 (38.5%) men. 21 women are represented in the training set (60%), 5 in the validation set (83%), and 6 in the test set (55%). The vast majority of participants report an ASL level of 6 or 7, as we show in Figure 5 in Appendix A. The participants also list their U.S. states. Using this information, we divide them into four regions based on the U.S. Census definitions<sup>4</sup>: Northeast, Midwest, South, and West. More participants in the dataset are from the Northeast than any other region, as shown in Figure 5 in Appendix A. We also find that the age range of participants is skewed; participants in their 20s and 30s make up 32 of the 52 participants (see Figure 6 in Appendix A).

Participants did not note their ethnicity or race for this dataset. As such, to uncover potential biases related to the participants’ perceived skin tone in their videos, we run the `skin-tone-classifier` Python package from Rejón Pina and Ma on the frame with the first detected face in each video. We find that when we do not specify that the videos were in color, the classifier most often detects them as black and white. When we specify that the videos are in color, the most common skin tone detected (out of the default color palette used in Rejón Pina and Ma) is `#81654f`. Because the classifier most commonly detects as black and white, we also try specifying the video frames as being black and white. In this setting, the most common skin tone detected is `#b0b0b0`, and the distribution differs from when the images are specified color images. We plot these results in Figure 7.

#### 3.2 Sign and Video Features

Because the ASL Citizen dataset is composed of signs from ASL-Lex (Caselli et al., 2017), we can utilize ASL-Lex’s lexical annotations for each sign. No works have studied these features in-depth on the ASL Citizen sign videos. We also analyze the video lengths, similarities and differences from the seed signer, and other video features.

**Video Length** We study the distribution of video lengths in order to better understand how video length may vary in this dataset. We find that the

<sup>4</sup>[https://www2.census.gov/geo/pdfs/maps-data/maps/reference/us\\_regdiv.pdf](https://www2.census.gov/geo/pdfs/maps-data/maps/reference/us_regdiv.pdf)

distribution of video lengths ( $s$ ) is skewed left, with a longer tail on the right, as shown in Figure 8.

We also study whether video lengths vary, on average, for participants of different ages and genders. To account for differences between the signs depicted by participants (since participants did not all record the same signs), for each video, we calculate the number of standard deviations (SDs) the video length is away from the mean for all videos of that sign - in other words, we calculate the  $z$ -score at the sign level. We show this calculation in the equation below, where  $v_i(s)$  represents the length of video  $i$  depicting sign  $s$ .

$$z = \frac{v_i(s) - \mu_s}{\sigma_s} \quad (1)$$

We find that, while men on average record videos over .3 SDs longer than the mean, women on average record videos over .2 SDs shorter than the mean. Thus, compared to other videos with the same sign, women record shorter videos than men on average. We show these results in Figure 9. Older participants, particularly those in their 70s, record longer videos on average (again, relative to other videos of the same sign) than younger participants. During manual inspection, we find older participants are more likely to have longer pauses before or after signing than younger participants, which may explain this gap. We also show these results in Figure 9.

**Sign Frequency** The ASL Citizen dataset is comprised of 2731 signs from the ASL-Lex dataset [Caselli et al. \(2017\)](#), a dataset with expert annotations about properties of each sign including frequency of use, iconicity, and varying phonological properties. To collect sign frequency labels, deaf signers who use ASL were asked to rate signs from 1 to 7 in terms of how often they appear in everyday conversations, where 1 was “very infrequently” and 7 was “very frequently”. We plot and compare the sign frequency distributions for the ASL Citizen dataset and the ASL-Lex dataset in Figure 10, and find that they are very similar.

We also find that there is little variation in average sign frequency for different genders. For male participants, the average sign frequency is 4.1592, while the average sign frequency for female participants is 4.1395, indicating that female participants chose slightly less frequently-occurring signs than men overall.

**Sign Iconicity** The ASL-Lex dataset also contains crowd-sourced annotations for sign iconicity,

where non-signing hearing annotators watch videos of a sign and evaluated how much they look like the sign’s meaning from 1 (not iconic) to 7 (very iconic). We calculate an average iconicity of 3.378 in the ASL-Lex dataset, and 3.379 in the ASL Citizen dataset. We plot these distributions in Figure 11, and again find that they are very similar.

We find average iconicity is 3.378 for women and 3.381 for men. This indicates that, as with frequency, there is only a slight difference, on average, between the iconicity of signs chosen by male and female participants.

## 4 Methods

Here, we describe the baselines for our ISLR experiments, along with the experimental settings we use.

### 4.1 Baselines

For our experiments, we use the baseline I3D and ST-GCN models which were trained on the ASL Citizen dataset and released along with the dataset.<sup>5</sup> We describe the details of these models below.

**I3D** The I3D model is a 3D convolutional network trained on the video frames themselves ([Carreira and Zisserman, 2017](#)). As with the original ASL Citizen baselines, we train our I3D model on preprocessed video frames from the sign videos in the ASL Citizen training set. These videos are each standardized to 64 frames by skipping or padding frames depending on video length. Videos are then randomly flipped horizontally to imitate right- and left-handed signers.

### 4.2 ST-GCN

The ST-GCN model is a temporal graph convolutional network trained on pose information ([Yan et al., 2018](#)). As with the original ASL Citizen baseline, we obtain pose representations for each frame using Mediapipe holistic ([Lugaresi et al., 2019](#)), with a set of 27 keypoints established by OpenHands ([Selvaraj et al., 2022](#)). These keypoints are center scaled and normalized using the distance between the shoulder keypoints. The frames are capped at a maximum of 128, and random shearing and rotation transformations are applied during training for data augmentation.

<sup>5</sup><https://github.com/microsoft/ASL-citizen-code>

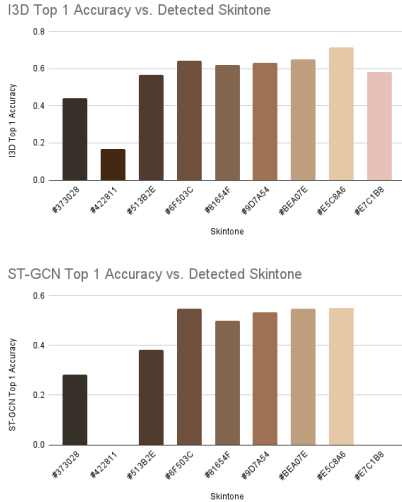


Figure 2: I3D (top) and ST-GCN (bottom) top 1 accuracy scores by detected skin tone. We find that, despite being less represented in the dataset, videos with lighter detected skin tones have higher accuracy scores on average for both models. The ST-GCN model, in particular, exhibits this behavior.

### 4.3 Experimental Settings

All baselines are run on a Mac Studio with an Apple M2 Max chip and 64GB RAM.

**I3D** We use the same experimental settings as the I3D ASL Citizen baseline: 75 epochs maximum, learning rate of  $1e-3$ , weight decay of  $1e-8$ , an Adam optimizer and ReduceLROnPlateau scheduler with patience 5. As described in the ASL Citizen paper, we calculate the loss by averaging cross-entropy loss and per-frame loss.

**ST-GCN** As with the original ASL-Citizen baseline, we train our ST-GCN model for a maximum of 75 epochs using a learning rate of  $1e-3$ , an Adam optimizer, and a Cosine Annealing scheduler.

## 5 Which factors impact dictionary retrieval results in the ASL Citizen dataset?

### 5.1 Participant-level differences

#### Baseline models perform over 10 percentage points better for male vs. female participants

We run the baseline I3D and ST-GCN models trained on the ASL Citizen dataset (Desai et al., 2024), and, for both models, find an accuracy disparity between male and female participants. For the I3D model, the overall Top-1 accuracy is 0.6306, while for females it is 0.5914 and for males it is 0.6776; in other words, a gap of over 10 points

in favor of male participants is observed. An even bigger gap is observed for the ST-GCN model; the overall Top-1 accuracy is 0.5944, while the Top-1 accuracy is 0.6838 for males and 0.52 for females.

**Average model accuracy varies greatly between participants** One possible contributor to the above performance disparities for male and female participants is variation in participant-level model accuracy. There are 11 participants whose videos are in the test set for the ASL Citizen dataset. Of these 11 participants, 6 are female and 5 are male. When calculating accuracy scores for each participant, we find high variation for both models, with over 15-point differences between the highest and lowest accuracy scores (see Table 5). This variation may contribute to the gender performance gap, as there are only a few participants of each gender in the test set.

While performing manual inspection, we find several characteristics of user videos that appear to vary between participants. Different participants have different background or lighting quality, and some participants mouth the word being signed while other participants do not. We also find instances of repetition, where the sign is repeated in the video, from P15, a female participant. There are also some instances of fingerspelling, where participants fingerspell the sign before signing it. These and other individual differences may contribute to the observed performance disparities.

#### The models perform better on lighter skin tones than darker skin tones on average

Despite darker skin tones making up most of the detected skin tones for videos in this dataset (see Figure 7), we find that models average higher performance when the detected skin tone is lighter. We illustrate this phenomenon for both models in Figure 2. As this figure shows, I3D follows similar trends to ST-GCN in terms of comparative performance for different skin colors, performing the best for lighter skin tones #BEA07E and #E5C8A6. That being said, ST-GCN performs comparatively more poorly on the three darkest skin tones (#373028, #422811, and #513B2E) and the lightest skin tone (#E7C1B8) than I3D, when compared to the higher-performing skin tones. This indicates that, though both models show similar patterns regarding the skin tones with higher/lower performances, the RGB-based I3D model appears to perform better overall on darker skintones than the ST-GCN model. Although we find variations in accuracy be-

Std. devs from mean	I3D Top-1	ST-GCN Top-1
$n < -2$	0.38462	0.3846
$-2 \leq n < -1$	0.5551	0.4862
$-1 \leq n < 0$	0.648	0.5888
$0 \leq n < 1$	<b>0.6704</b>	<b>0.6449</b>
$1 \leq n < 2$	0.5727	0.5878
$n > 2$	0.3846	0.4668

Table 1: Top-1 accuracy scores for videos within a certain number of SDs away from the mean for videos of the same sign. For both models, videos with lengths closer to the mean yield better model performance.

tween participants in the previous section, the skin tones are categorized at the video level. Thus, it is possible to see variation in predicted skin tone for different videos recorded by the same individual. The lighting quality of individual videos may be a confounder for these results.

### Trained models exhibit the highest average performance on participants in their 20s and 60s

The ASL Citizen test set is made up of 11 individuals in their 20s, 30s, 50s, and 60s. We find that, as with gender, model accuracy varies for different age ranges; the highest accuracy scores were achieved for participants in their 20s and 60s. This could be influenced by the proportion of participants in their 20s in the training set.

## 5.2 Video-level differences

### Performance decreases as the video length diverges from the average

For each sign video in the ASL Citizen dataset, we calculate the  $z$ -score of its video length compared to other videos of the same sign. We then place these values into buckets: less than -2, -2 to -1, -1 to 0, 0 to 1, 1 to 2, and more than 2 SDs from the mean. We find that, on average, the videos farther away from the mean see decreased model performance compared to the videos closest to the mean. The results in full are in Table 1.

### Performance decreases when video quality degrades

In addition to video length, we study the impact of video quality on model accuracy. Given that we are studying the quality of individual video frames without a reference image, we use the BRISQUE score (Mittal et al., 2012) to measure image quality of individual frames. Higher BRISQUE scores indicate lower quality, while lower BRISQUE scores indicate higher quality. We find that higher BRISQUE scores correlate negatively with Top-1 model performance for




Figure 3: Association between BRISQUE image quality scores and accuracy. Higher BRISQUE scores indicate lower image quality, and vice versa. Thus, higher image quality appears to be associated with better model performance.

the I3D model, with a Spearman correlation of  $\rho = -0.0367$  and a  $p$ -value of  $p = 1.53 \times 10^{-8}$ . We show a scatterplot of these results in Figure 3, along with a linear regression line.

### Dissimilarity between participant and seed signer signs negatively impacts model accuracy for the ST-GCN pose model

The Fréchet distance is often used as an evaluation metric for sign language generation, to study the similarity between generated signs and references (Hwang et al., 2024; Dong et al., 2024) (see § D for more details). In the ASL Citizen dataset, one of the participants is a paid ASL model who records videos for every sign, referred to as the “seed signer”.

We study whether dissimilarity between the participant and seed signer may have a negative impact on model accuracy. To do so, we use the pose models used as input to the ST-GCN model. Every .25 seconds, we measure the distance between the model pose and the participant’s pose at that frame, studying the distance between left hands and right hands separately. We find no significant relationship between right hand or left hand distance from the seed signer for the I3D model, and for the ST-GCN model we find a significant negative Spearman correlation between distance from the seed signer and accuracy for the right hand ( $\rho = -.0289$ ,  $p = 0.001$ ). We plot these results, along with lines of best fit, in Figure 12.

### When the average signing “speed” is closer to the sign-level average, performance is better

In addition to video length, we are interested in studying the average distance between poses over consistent time intervals. We want to study how much movement on average occurs within these increments, i.e. the “speed” of sign production. We study this by calculating the pairwise Fréchet distance between poses at each 0.25 second interval,

SD from mean	I3D (LH)	ST-GCN (LH)	I3D (RH)	ST-GCN (RH)
$n < -2$	.4627	.2139	.5	.2375
$-2 \leq n < -1$	.6041	.5804	.6121	.5174
$-1 \leq n < 0$	.6503	.6426	.6438	.6351
$0 \leq n < 1$	.6244	.5813	.6423	.6145
$1 \leq n < 2$	.6164	.5261	.616	.5744
$n > 2$	.5711	0.4739	.5619	.5107

Table 2: Number of SDs away from the mean of the sign (in buckets) for the “speed” of signing, i.e. the average Frechet distance between poses every 0.25 seconds, for right hand and left hand. We find that, for both right hand and left hand, the performance degrades as the average “speed” of the sign production in a sign video deviates from the average for that particular sign.

with distance calculated between a pose and the pose .25s after, starting from the first frame. We again take this distance for the participants’ right hand and left hand. We find that, on average, the farther away a participant’s average signing speed is from that sign’s mean, the worse performance is, with especially high performance degradations 2 SDs or more from the mean. We show these results in Table 2.

### 5.3 Sign-level lexical features

Here, we present results for four sign-level features annotated in the ASL-Lex dataset: sign frequency, iconicity, phonological complexity, and neighborhood density. We find that several of these features are significantly correlated with model performance, which we discuss below.

**Sign frequency, phonological complexity, and neighborhood density are negatively correlated with model accuracy** As mentioned in § 3.2, sign frequency annotations in the ASL-Lex dataset were collected from ASL signers. The ASL-Lex 2.0 dataset (Sehyr et al., 2021) also contains a new phonological complexity metric. Using 7 different categories of complexity, scores were calculated by assigning a 0 or 1 to each category (depending on whether that category was present) and adding them together, for a maximum possible scores of 7 (most complex) and a minimum possible score of 0. The highest complexity score in the dataset was a 6. Neighborhood density was calculated based on the number of signs that shared all, or all but one, phonological features with the sign.

Intuitively, we expect negative associations with phonological complexity and accuracy as well as neighborhood density and accuracy, and in-

deed find significant negative correlations ( $\rho = -0.0618$ ,  $p = 0.005$  for phonological complexity and  $\rho = -0.0584$ ,  $p = 0.01$  for neighborhood density). However, we also find a significant negative association between sign frequency and model accuracy ( $\rho = -0.057$ ,  $p = 0.011$ ). Existing work indicates that higher-frequency words are produced more quickly than low-frequency words (Jescheniak and Levelt, 1994; Emmorey et al., 2013; Gimeno-Martínez and Baus, 2022); thus, it is possible that this association could be related to video length.

**There is no significant correlation between iconicity and model accuracy** As mentioned in § 3.2, sign iconicity ratings were also collected for the ASL-Lex dataset. We find a very slight positive correlation between sign iconicity and model accuracy ( $\rho = 0.044$ ), which is not significant ( $p = 0.8424$ ). Thus, we conclude that visual similarity to the English word appears not to affect the model’s ability to recognize a sign.

### 5.4 Which features are the best predictors of model accuracy?

After looking at the impacts of lexical, demographic, and video features on model accuracy, we are interested in studying which features are (by themselves) the best predictors of model accuracy. As such, we study the mutual information between each feature and the Top-1 accuracy for the I3D and ST-GCN models. We study 19 features in total, where some relate to participant demographics (e.g. age and gender), others relate to the sign lexical features (e.g. sign iconicity), and the rest are characteristics of individual videos (e.g. BRISQUE score and Frechet distances). We find that the 5 most impactful features are characteristics of individual videos (BRISQUE, Frechet from seed signer, and absolute  $z$ -score of “signing speed”), with BRISQUE video quality scores showing the highest mutual information with Top-1 accuracy. Out of the lexical features, sign iconicity has the highest mutual information, and out of the demographic features, the participant’s ASL level has the highest mutual information with the model performance. The results are in Table 6.




Figure 4: The relationships between sign frequency (left), sign iconicity (center left), phonological complexity (center right), and neighborhood density (right) and top 1 accuracy for the ST-GCN model. We find that sign frequency, phonological complexity, and neighborhood density are all significantly negatively correlated with model accuracy ( $p < 0.05$ ) when calculating the Spearman’s rank correlation. However, despite a slight positive correlation between iconicity and accuracy, the  $p$ -value is not significant.

Model	Overall			Female participants			Male participants			Parity (Top-1)
	Top-1	Top-5	Top-10	Top-1	Top-5	Top-10	Top-1	Top-5	Top-10	
ST-GCN	.5238	.7665	.8295	.4406	.6886	.7665	.6236	.8601	<b>.9374</b>	.7065
ST-GCN (VL)	.5488	.7923	.8515	.4666	.7200	.7941	.6476	.8791	.9205	.7205
ST-GCN (VL, fem.)	.5395	.7926	.8538	.4621	.7202	.7974	.63	<b>.8795</b>	.9216	.7334
ST-GCN (brisque, HP)	.4723	.7344	.8046	.3949	.6551	.7354	0.5653	.8296	.8877	.6986
ST-GCN (brisque, LP)	<b>.5580</b>	<b>.7960</b>	<b>.8545</b>	<b>.4801</b>	<b>.7279</b>	<b>.8011</b>	<b>.6516</b>	.8779	.9187	<b>.7368</b>

Table 3: Performance of ST-GCN baseline against models that use the resampling strategies discussed in 6.3. We find that all resampling strategies improve accuracy and gender parity over the baseline (for every metric but Top-10 Male), and resampling lower quality videos at a higher rate improves gender parity the most, followed closely by resampling based on video length from only female participants.

## 6 Can we mitigate disparate impacts while maintaining high model performance for dictionary retrieval?

### 6.1 Training on single-gender subsets

We first address the gender performance gap by training on participants of each gender in isolation. When doing this, we find a slight difference between the performance gaps for models trained on male-only and female-only subsets. For the model trained on the male-only subset, the Top-1 accuracy for male subjects is .292, and the Top-1 accuracy is .168. For the model trained on the female-only subset, the Top-1 accuracy for male subjects is .291, and the Top-1 accuracy for female subjects is .206. Thus, the model trained only on female subjects has a smaller gap, and higher accuracy parity, between male and female subjects than the model trained on only male subjects. However, both models have low performance overall, so the Top-1 accuracy parity for subjects of different genders (calculated by dividing the female accuracy by the male accuracy) is .7571 for the model trained on all subjects compared to .7079 for the model trained on only female subjects. The model trained on only male subjects has the lowest accuracy parity, at .5746. We show these results in Table 7 in Appendix I.

### 6.2 Training label shift

In addition to training on single-gender subsets, we experiment with a label-shift approach to debiasing. Because ISLR is a multiclass problem, we experiment with the reduction-to-binary approach for debiasing multi-class classification tasks proposed by Alabdulmohsin et al. (2022). We run the label-shift algorithm and train the ST-GCN model on the debiased labels for 25 epochs, and compare the performance of the debiased model to the ST-GCN model without debiasing, which we also train for 25 epochs. We find that the model trained on regular labels actually has a *higher* ratio for female to male accuracy than the debiased model: .7476 for the baseline model, and .7052 for the debiased model. We show these results in full in Table 8.

### 6.3 Weighted resampling

Although there is a large gender performance gap observed (§5.1), based on the results from Table 6, other features are much more heavily tied to model accuracy. Thus, it is likely that these features (in particular, features at the video level) may influence results. But what happens if the impact of videos with potentially-noisy features is reduced during training? We experiment with weighted resampling, where samples with certain features are



more likely to be resampled. We explain how we calculate the resampling probability, and present results, for each variable we study in the paragraphs below.

**Video length** We first experiment with calculating the resampling probability based on video length. Given that videos closer to the mean produced higher accuracy scores, we wanted to resample these videos at a higher rate to reduce training noise. We calculate the probability of resampling as follows, where  $v_i(s)$  refers to the length of video  $i$  for sign  $s$ ,  $\mu_s$  refers to the mean video length of videos depicting sign  $s$ , and  $\sigma_s$  refers to the SD for video lengths of videos depicting sign  $s$ :

$$P(\text{resample}) = \frac{1}{2^{\frac{v_i(s) - \mu_s}{\sigma_s}}} \quad (2)$$

We show the results for this approach in Table 3, represented by the ST-GCN (VL) model. We find that this approach improves upon the baseline ST-GCN model by at least 2 percentage points for all accuracy metrics, and improves gender parity for Top-1 accuracy by 1.4%.

**Video length for female participants** We then experiment with the exact same resampling process described above, based on number of SDs from the mean for video length, but only resample videos from female participants. Because training on an all-female subset yielded a higher test accuracy for female subjects than an all-male subset (Table 7), we want to investigate whether restricting our resampled data to female participants improves the gender performance gap. We show these results in Table 3, under the baseline ST-GCN (VL, fem.). We find that this approach exceeds calculating the resampling probability using video length for participants of all genders for Top-5 and Top-10 accuracy. We also find that this baseline achieves the second-highest gender parity of all of the baselines, at 2.69% higher than the baseline. Thus, we find evidence that resampling based on video length SDs, but only videos from female participants (the group with the lower model accuracy scores), greatly improves gender parity over the baseline model.

**BRISQUE score** Because the BRISQUE score shows the highest mutual information with Top-1 accuracy, we experiment with resampling based on the video quality. We experiment with two different resampling strategies: resampling higher-quality videos at a higher rate (*resampling high quality*)

and resampling lower-quality videos at a higher rate (*resampling low quality*). We discuss these strategies below.

**Resample high quality:** We first experiment with resampling more high-quality videos (lower BRISQUE scores) at a higher rate by setting the resampling probability as a function of the BRISQUE score, with higher BRISQUE scores reducing the resampling probability. We calculate the probability of resampling as follows, where  $B_i(s)$  refers to the BRISQUE score of video  $i$ :

$$P(\text{resample}) = \frac{1}{2^{\frac{B_i}{100}}} \quad (3)$$

**Resample low quality:** We then experiment with resampling more low-quality videos (higher BRISQUE scores) at a higher rate by setting the resampling probability as a function that increases relative to the BRISQUE score. We calculate the probability of resampling as follows, where  $B_i(s)$  refers to the BRISQUE score of video  $i$ :

$$P(\text{resample}) = \frac{1}{2^{\frac{100}{B_i}}} \quad (4)$$

Our results in Table 3 show that the latter approach, *resample low quality*, achieves the highest overall accuracy and gender parity score.

## 7 Conclusion

In this work, we address a gap in sign language processing research by exploring biases in sign language resources, and experimenting with strategies to mitigate these biases. We focus on the ASL Citizen dataset in particular, and release demographic information for this dataset to aid future work. We find performance gaps related to skin tone, participant age, and gender. Still, we find that video level features, such as the video quality, signing "speed", and video length, appear to be the best predictors of model accuracy. We find that selectively resampling data with video lengths closer to the mean improves overall performance. We also find that doing this resampling strategy for *only* the group with lower model performance (female, when comparing genders) improves the gender parity for model performance. We find that resampling lower-quality videos at a higher rate achieves the highest Top-1 accuracy *and* gender parity.

## Limitations

While in this work we find and document performance gaps between participants of different demographics such as age and gender, because of

the differences between individual participants that we detail above (see Table 5), and the number of participants in the test set (11), it is unclear how much of these differences are due to age or to other underlying factors.

Another limitation is that we focus on a single dataset. This is due in part to the fact that this is the only large-scale crowdsourced dataset for isolated sign language recognition with demographic labels. However, as more crowdsourced sign language resources become available, it is critical that these analyses are repeated on these datasets to assess the generalizability of our results.

## Ethical Implications

In our analysis of participant demographics, and accompanying features, for the ASL Citizen dataset, we present some characteristics of the dataset that vary between demographics. For instance, we discuss our findings that male participants and older participants typically record longer videos. It is important to emphasize that these findings should not be generalized to all ASL signers, and that they should instead be used to study the characteristics of this dataset in particular.

Further, this work is not exhaustive; there are many sources of bias unexplored by this work, including differences in participant culture or ethnicity. There may be many more sources or dimensions of bias not covered in this paper that should be explored by future work.

We also note that participants who chose to denote their demographic information (which was optional) consented for this information to be anonymously released as part of the dataset. No identifiable information about the participants will be released with the publication of this paper; rather, anonymous participant IDs will be accompanied with their demographics.

## Acknowledgments

We would like to thank all of the participants who contributed videos to the ASL Citizen dataset, without whom this work would not have been possible.

## References

Ibrahim Mansour I Alabdulmohsin, Jessica Schrouff, and Sanmi Koyejo. 2022. [A reduction to binary approach for debiasing multiclass datasets](#). In *NeurIPS 2022*.

Vassilis Athitsos, Carol Neidle, Stan Sclaroff, Joan Nash, Alexandra Stefan, Quan Yuan, and Ashwin Thangali. 2008. The american sign language lexicon video dataset. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 1–8. IEEE.

Carole Sue Bailey, Kathy Dolby, and Hilda Marian Campbell. 2002. *The Canadian dictionary of ASL*. University of Alberta.

Patrick Boudreault and Rachel I Mayberry. 2006. Grammatical processing in american sign language: Age of first-language acquisition effects in relation to syntactic structure. *Language and cognitive processes*, 21(5):608–635.

Joao Carreira and Andrew Zisserman. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308.

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*, 49:784–801.

Kearsy Cormier, Adam Schembri, and Bencie Woll. 2010. Diversity across sign languages and spoken languages: Implications for language universals. *Lingua*, 120(12):2664–2667.

Aashaka Desai, Lauren Berger, Fyodor Minakov, Nessa Milano, Chinmay Singh, Kriston Pumphrey, Richard Ladner, Hal Daumé III, Alex X Lu, Naomi Caselli, et al. 2024. Asl citizen: A community-sourced dataset for advancing isolated sign language recognition. *Advances in Neural Information Processing Systems*, 36.

Lu Dong, Lipisha Chaudhary, Fei Xu, Xiao Wang, Mason Lary, and Ifeoma Nwogu. 2024. [Signavatar: Sign language 3d motion reconstruction and generation](#). Preprint, arXiv:2405.07974.

Thomas Eiter, Heikki Mannila, and Christian Doppler Labor für Expertensysteme. 1994. Computing discrete fréchet distance.

Karen Emmorey, Ursula Bellugi, Angela Friederici, and Petra Horn. 1995. Effects of age of acquisition on grammatical sensitivity: Evidence from on-line and off-line tasks. *Applied psycholinguistics*, 16(1):1–23.

Karen Emmorey, Jennifer AF Petrich, and Tamar H Golman. 2013. Bimodal bilingualism and the frequency-lag hypothesis. *Journal of deaf studies and deaf education*, 18(1):1–11.

Marc Gimeno-Martínez and Cristina Baus. 2022. [Iconicity in sign language production: Task matters](#). *Neuropsychologia*, 167:108166.

- Eui Jun Hwang, Huije Lee, and Jong C. Park. 2024. [Autoregressive sign language production: A gloss-free approach with discrete representations](#). *Preprint*, arXiv:2309.12179.
- Jörg D Jescheniak and Willem JM Levelt. 1994. Word frequency effects in speech production: Retrieval of syntactic information and of phonological form. *Journal of experimental psychology: learning, Memory, and cognition*, 20(4):824.
- Lee Kezar, Jesse Thomason, Naomi Caselli, Zed Sehyr, and Elana Pontecorvo. 2023a. [The sem-lex benchmark: Modeling asl signs and their phonemes](#). In *Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility, ASSETS '23*, New York, NY, USA. Association for Computing Machinery.
- Lee Kezar, Jesse Thomason, and Zed Sehyr. 2023b. [Improving sign recognition with phonology](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2732–2737, Dubrovnik, Croatia. Association for Computational Linguistics.
- Annelies Kusters. 2010. Deaf utopias? reviewing the sociocultural literature on the world’s “martha’s vineyard situations”. *Journal of deaf studies and deaf education*, 15(1):3–16.
- Dongxu Li, Cristian Rodriguez, Xin Yu, and Hongdong Li. 2020. Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1459–1469.
- Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. 2019. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*.
- Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. 2012. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 21(12):4695–4708.
- Carol Padden. 2010. Sign language geography. *Deaf around the world: The impact of language*, pages 19–37.
- René Alejandro Rejón Pina and Chenglong Ma. [Classification algorithm for skin color \(casco\): A new tool to measure skin color in social science research](#). *Social Science Quarterly*, n/a(n/a).
- Noha Sarhan, Christian Wilms, Vanessa Closius, Ulf Brefeld, and Simone Frintrop. 2023. [Hands in focus: Sign language recognition via top-down attention](#).
- 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.
- Prem Selvaraj, Gokul Nc, Pratyush Kumar, and Mitesh Khapra. 2022. [OpenHands: Making sign language recognition accessible with pose-based pretrained models across languages](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2114–2133, Dublin, Ireland. Association for Computational Linguistics.
- William C Stokoe, Dorothy C Casterline, and Carl G Croneberg. 1965. *A dictionary of American Sign Language on linguistic principles*. Gallaudet College Press, Washington, DC.
- Andrea Toliver-Smith and Betholyn Gentry. 2017. Investigating black asl: A systematic review. *American Annals of the Deaf*, 161(5):560–570.
- Ronnie Wilbur and Avinash C Kak. 2006. Purdue rvl-slll american sign language database.
- Sijie Yan, Yuanjun Xiong, and Dahua Lin. 2018. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Morteza Zahedi, Daniel Keysers, Thomas Deselaers, and Hermann Ney. 2005. Combination of tangent distance and an image distortion model for appearance-based sign language recognition. In *Pattern Recognition: 27th DAGM Symposium, Vienna, Austria, August 31-September 2, 2005. Proceedings 27*, pages 401–408. Springer.




Figure 5: Distribution of ASL levels (left) and regions (right) of participants for the ASL Citizen dataset.




Figure 6: Age ranges of participants in the ASL Citizen dataset. Participants are skewed mostly towards their 20s and 30s, with a lesser skew towards participants in their 60s.

## A Participant Demographics

Here, we plot the demographic information discussed in 3.1. Note that providing demographic information was optional, so these numbers will not always add up to the total number of participants (52).

In Figure 5, we plot the distribution of ASL levels and regions associated with the participants in the ASL Citizen dataset. We find that most participants are at an ASL level of 6 or 7, with only one participant each at level 3 or 4. A plurality of participants are from the Northeast, almost half. The West contains the fewest participants.

In Figure 6, we plot the distribution of participants' ages. We find that participants are mostly skewed towards younger adults (20s and 30s) but that there is also a slight skew towards contestants in their 60s. Contestants in their 20s, 30s, 40s, 50s, 60s, and 70s are represented in the dataset, but contestants in their 40s and 70s are not represented in the test set.

In Figure 7, we plot the distribution of skin tones in the dataset when frames are set as color images and black-and-white images. We include black-and-white images because we found that, when an image type was not set, the model detected the





Figure 7: Frequency of detected skin tones of participants in videos when the video frames were set manually to color images (left) and black and white images (right)

images as black-and-white images in the majority of cases. One notable finding is that the skin color model detected lighter skin tones more frequently when the images were set to black-and-white than when they were set to color images. This indicates possible unreliability of the skin color detection; it is possible, for instance, that when the images are set to color, the system classifies the skin colors as darker than they actually are.

## B Video Length Distributions

In Figure 8, we find that video lengths have a skewed distribution, where the average video length is higher than the median. In other words, video lengths lower than the mean are more common and vice versa, and there is a long tail to the right. After watching participants' videos, we suspect that this difference in video length is a result of some participants having a tendency to pause for multiple seconds at the beginning or end of their recording. This happens especially often with the first couple of videos that people record.

We also find that female participants have, on average, shorter videos related to their signs than male participants. For each sign video, we calculated the mean and standard deviation for all videos with that sign. We then calculated how many standard deviations those movies were away from the mean.




Figure 8: Distribution of video lengths for all sign videos in the ASL Citizen dataset. The distribution is skewed towards the right, with a long tail on the right.

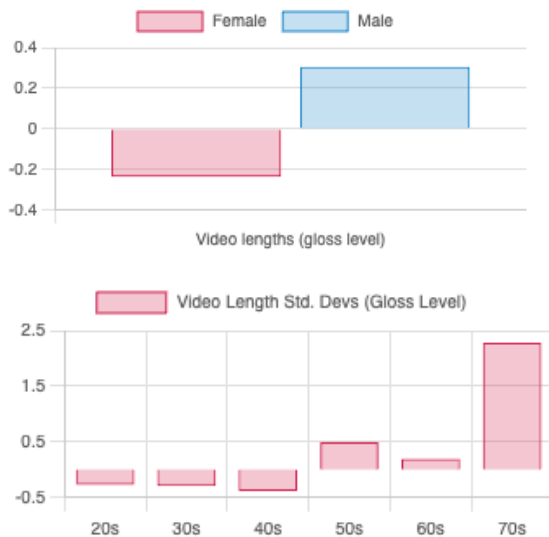


Figure 9: Average number of standard deviations away from the mean at the sign level for male and female participants (top) and participants in their 20s, 30s, 40s, 50s, 60s, and 70s (bottom). Relative to other videos of the same sign, women tend to record shorter videos, and older participants tend to record longer videos.

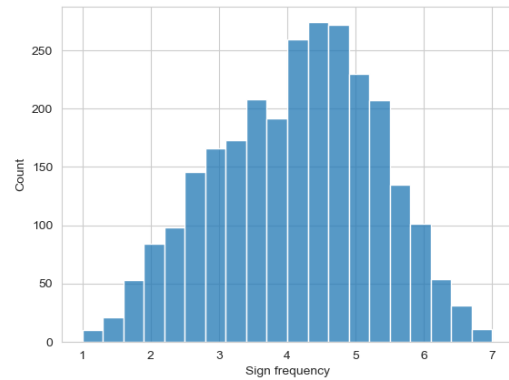


Figure 10: Distributions of labeled sign frequencies for each of the 2731 signs from the ASL-Lex dataset (top) and all of the sign videos in the ASL Citizen dataset (bottom). The distributions are very similar, indicating that users chosen signs of certain frequencies at a similar rate to how they are distributed in the ASL-Lex dataset.

### C Lexical Feature Distribution

In addition to getting demographic and video features, we used the ASL-Lex (Caselli et al., 2017) annotations to analyze lexical features in the ASL Citizen dataset. We found that, for sign frequency and iconicity, the distributions are very similar to those in the ASL-Lex dataset. The distributions of both datasets are plotted side-by-side for frequency and iconicity, respectively, in Figures 10 and 11.

### D Frechét Distance

The Frechét distance, used as a similarity metric between curves, and is commonly described in the following manner:

A man is walking a dog on a leash: the man can move on one curve, the dog on the other; both may vary their speed, but backtracking is not allowed. What is the length of the shortest leash that is


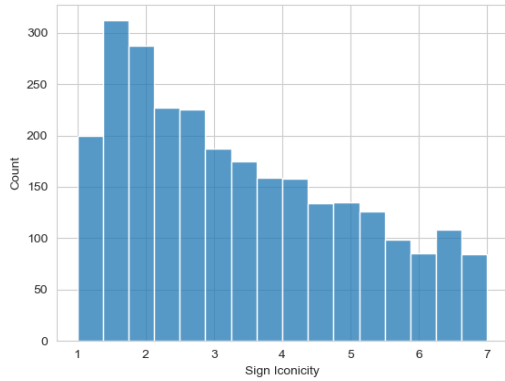


Figure 11: Distribution of sign iconicities in the ASL-Lex dataset (left) and the sign videos recorded in the ASL Citizen dataset (right). Like the sign frequencies, the iconicities in the ASL Citizen videos are distributed similarly to their distribution in the ASL-Lex dataset.

Age range	# in test	I3D Top-1	ST-GCN Top-1
20s	2	.6697	.6076
30s	3	.5689	.5336
40s	0	–	–
50s	2	.549	.5658
60s	3	<b>.7016</b>	<b>.6421</b>
70s	0	–	–

Table 4: Average accuracy scores for participants of each age range in the test set. There were no participants in their 40s or 70s in the test set, and one participant did not specify their age. We find the highest performance in both models occurs for participants in their 20s and 60s.

Participant ID	I3D Top-1	ST-GCN Top-1
<b>P6</b>	0.5456	0.4387
<b>P9</b>	0.6586	0.5663
<b>P15</b>	0.4653	0.5757
<b>P17</b>	0.6183	0.4997
<b>P18</b>	0.7065	0.5727
<b>P22</b>	0.5562	0.4671
<b>P35</b>	0.7204	0.7153
<b>P42</b>	0.6041	0.6949
<b>P47</b>	0.7471	0.7886
<b>P48</b>	0.6882	0.6652
<b>P49</b>	0.6327	0.556

Table 5: Model top-1 accuracy scores on the set of videos recorded by each participant in the test set. For both models, there is high variation between participants, with scores ranging from 0.4653 to 0.7204 (I3D) and 0.4387 to 0.7886 (ST-GCN).

sufficient for traversing both curves?  
- (Eiter et al., 1994)

## E Accuracies for different age ranges

In Table 4, we show the Top-1 accuracy scores for the I3D and ST-GCN model for participants of different ages. We find the highest scores occur for participants in their 20s and 30s, with the third highest scores occurring for participants in their 60s. Participants in their 40s and 70s were not represented in the test set.

## F Model accuracies for each participant in the test set

In Table 5, we report the accuracy scores for the baseline ST-GCN model on the participants in the test set of the ASL Citizen dataset. We find differences of over 20 points between participant averages for both models. P6, P9, P15, P17, P18, and P22 disclosed that they are female, while the other participants disclosed that they are male.

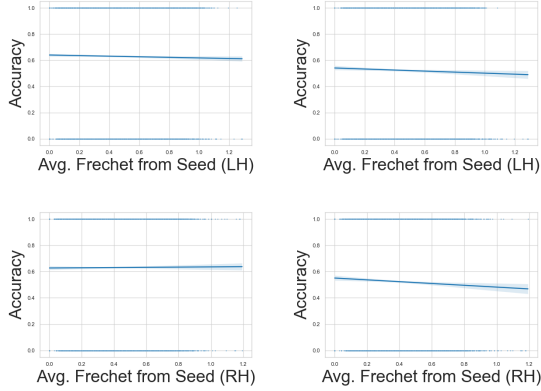


Figure 12: The Frechet distance from the seed (model) signer vs. top-1 accuracy for the I3D model (top) and ST-GCN model (bottom), with the distance between left hands on the left and the distance between right hands on the right.

## G Frechet distance from seed signer

In Figure 12, we plot the Top-1 accuracies for the I3D and ST-GCN model as a function of the Frechet distance from the seed signer for each sign video (where the seed signer is a recruited ASL model for the ASL Citizen dataset). We find a significant negative correlation between Frechet distance from the seed signer and Top-1 accuracy for the ST-GCN pose model, but no significant correlations for the I3D model.

## H Mutual Information Results

In Table 6, we present the mutual information results in full for each studied variable. We study 19 variables total, spanning demographics, sign lexical features, and video-level features, and calculate the mutual information between each feature and the Top-1 accuracy. We find the highest levels of mutual information to occur for video-level features, suggesting features of individual videos are more impactful for model accuracy than demographic characteristics of the participants. Out of the demographic characteristics, the ASL level of the participant appears to be the most influential with respect to accuracy.

## I Results for models trained on single-gender subsets

Here, we report the model results for the ST-GCN model trained on single-gender subsets, comparing models trained on all-male and all-female subsets to the model trained on all of the training data. In

Feature	Mut. Info (ST-GCN)	Mut. Info (I3D)
BRISQUE	0.6920	0.6617
Avg. Frechet from seed (RH)	0.6444	0.6217
Abs. Avg. Frechet SD (RH)	0.6390	0.6090
Abs. avg. Frechet SD (LH)	0.6285	0.5641
Avg. Frechet from seed (RH)	0.5889	0.5403
Sign Iconicity	0.0757	0.0508
Sign Frequency	0.0619	0.0440
Abs. avg. Video Length SD	0.0293	0.0399
ASL Level	0.0048	0.0020
Region	0.0034	0.0002
Neighborhood Density	0.0032	0.0026
Number Of Morphemes	0.0026	0.0012
Phonological Complexity	0.0013	0.0006
Lexical Class	0.0007	0.0008
Iconicity Type	0.0002	0.0002
Gender	0	0.0034
Age	0	0.01107
Bounding Box Area (RH)	0	0
Bounding Box Area (LH)	0	0

Table 6: Mutual information for each of the features above and the Top-1 accuracy for the ST-GCN and I3D models, respectively. For both models, the BRISQUE score, average Frechet distance from the model (right hand and left hand) and the absolute value of the number of SDs of the average Frechet distance between frames are the top three features, with the other features far behind. This seemingly indicates that video-level features are the biggest indicator of model accuracy.

Table 7, we report the Top-1, Top-5, and Top-10 accuracy scores for each model.

## J Results for model trained on debiased labels

We report the results for a model trained for 25 epochs on training labels that were debiased using the reduction-to-binary techniques proposed by Alabdulmohsin et al. (2022). We find that the model trained on regular labels actually had a higher accuracy parity score (ratio of female accuracy to male accuracy) than the model trained on debiased labels. We show the Top-1, Top-5, and Top-10 results for each model in Table 8.

	Trained on female subjects			Trained on male subjects			Trained on all subjects		
	Top-1	Top-5	Top-10	Top-1	Top-5	Top-10	Top-1	Top-5	Top-10
All	.244	.479	.581	.224	.434	.527	.594	.828	.881
Male	.291	.548	.653	.292	.538	.639	.684	.902	.939
Female	.206	.421	.521	.168	.347	.433	.520	.767	.833

Table 7: Performances for ST-GCN model trained on only male subjects, only female subjects, and all subjects, respectively. We find that the model trained on only female subjects has the lowest performance gap between male and female subjects in the test set, but the ratio of female accuracy to male accuracy is highest for the model trained on all subjects.

	ST-GCN			ST-GCN (debiased)		
	Top-1	Top-5	Top-10	Top-1	Top-5	Top-10
All	.5323	.7997	.8622	.4821	.7576	.8265
Male	.6173	.8781	.9254	.5746	.8493	.9014
Female	.4615	.7343	.8096	.4052	.6811	.7641

Table 8: Performances for ST-GCN model trained on regular training labels (left) and debiased training labels (right). We find that the accuracy parity, calculated as the ratio of female to male accuracy, is higher for the model trained on regular training labels than the debiased model.