

Recognizing American Sign Language Nonmanual Signal Grammar Errors in Continuous Videos

Elahe Vahdani, Longlong Jing, and Yingli Tian*

The Graduate Center and The City College,
The City University of New York, New York, NY, USA.
{evahdani, ljing}@gradcenter.cuny.edu
ytian@ccny.cuny.edu (*Corresponding author)

Matt Huenerfauth

Golisano College of Computing and Information Sciences
The Rochester Institute of Technology
Rochester, NY, USA.
matt.huenerfauth@rit.edu

Abstract—As part of the development of an educational tool that can help students achieve fluency in American Sign Language (ASL) through independent and interactive practice with immediate feedback, this paper introduces a near real-time system to recognize grammatical errors in continuous signing videos without necessarily identifying the entire sequence of signs. Our system automatically recognizes if a performance of ASL sentences contains grammatical errors made by ASL students. We first recognize the ASL grammatical elements including both manual gestures and nonmanual signals independently from multiple modalities (i.e. hand gestures, facial expressions, and head movements) by 3D-ResNet networks. Then the temporal boundaries of grammatical elements from different modalities are examined to detect ASL grammatical mistakes by using a sliding window-based approach. We have collected a dataset of continuous sign language, ASL-HW-RGBD, covering different aspects of ASL grammars for training and testing. Our system is able to recognize grammatical elements on ASL-HW-RGBD from manual gestures, facial expressions, and head movements and successfully detect 8 ASL grammatical mistakes.

Keywords – American Sign Language; Grammar Recognition; Immediate Feedback; Multimodality; Deaf and Hard of Hearing.

I. INTRODUCTION

A. Motivation and Challenges

American Sign Language (ASL) has become the 4th most studied language at U.S. colleges [1]. To assist ASL students practicing their signing skills, we design a near real-time deep learning-based system to process continuous signing videos and evaluate students' performance. Our system is able to recognize a set of ASL grammatical mistakes by processing multiple modalities including hand gestures, facial expressions, and head movements from continuous videos.

Compared to isolated sign language recognition [2], continuous sign language recognition (CSLR) is more complicated because no temporal boundaries of signs are provided. Moreover, the transition movements between two consecutive signs are subtle and diverse and therefore hard to detect. Our goal is to recognize grammatically important components and detect possible mistakes without fully recognizing each individual word in continuous ASL videos.

ASL is a natural language with a distinct grammar structure from English, as illustrated by how questions may be formed in each language: In English, “*WH-Questions*” use words like



Fig. 1. Examples of ASL grammatical elements from multimodalities including facial expressions, head movements, and hand gestures: WH-Question (WHQ), Yes-No-Question (YNQ), Negative (NEG), Conditional (COND), Topic, and Rhetorical Question (RHQ).

Who, What, Where, When, Why, typically at the beginning of a sentence, e.g. “What did she buy yesterday?” In ASL, WH-word may occur in other positions in the sentence, including at the end, e.g. “SHE BUY YESTERDAY WHAT.” The question is primarily indicated by a “nonmanual signal” during a WH-word or a longer span of the sentence. Nonmanual signals may comprise several body movements such as head, shoulders, torso, eyebrows, eyeballs (gaze), eyelids, nose, mouth, tongue, cheeks, and chin [3]. The nonmanual signal for WH-Questions consists of a furrowing of brows and a tilt of head forward. In English, to change a declarative sentence to a “*Yes-No-Question*”, one can change the order of words from “She was there” to “Was she there?”. In ASL, to express a “*Yes-No-Question*”, rather than changing the word order, the question is indicated by a nonmanual signal (eyebrows raised and head tilted forward), especially at the end of the sentence [4].

Nonmanual signals are used to convey a variety of grammatical information in ASL, e.g. Topicalization, “*Yes-No-Question*”, “*WH-Question*”, etc. Because they consist of movements of multiple body parts, our recognition framework is built upon multiple modalities to detect these ASL grammatical elements. Note that different nonmanual signals may share behavioral properties, increasing the recognition difficulty. For example, during “*WH-Question*” and “*Yes-No-*

TABLE I
CLASSES OF MANUAL SIGNS DETECTED BY OUR SYSTEM.

Gesture Type	Description and Instances
Conditional	IF-SUPPOSE
Negative	NEVER, NONE, NOT, DON'T MIND
YNQ	QUESTION MARK WIGGLE, QUESTION
WHQ	WHO, WHAT, WHERE, WHY, WHEN
Time	NOW, TODAY, ALWAYS, LAST-WEEK
Pointing	IX-HE-SHE-IT, I-ME, YOU
Fingerspelling	Spelling out the letters by handshapes.
Clause Boundary	Gap between two consecutive clauses.

Question”, head tilts forward for both while eyebrows are furrowed during the former and raised during the latter, see Fig. 1. To address this challenge, we extract fine-grained features from different modalities.

B. ASL Linguistic Elements Recognized by Our System

A grammatically correct ASL sentence requires signer to synchronize manual signs (primarily consisting of gestures of the hands, see Table I) and nonmanual signals (primarily consisting of facial expressions and head movements, see Table II). For instance, if someone performs a WH-question sentence in ASL without performing the WH-Question nonmanual signal during the WH-word, this would be ungrammatical. As another example, if someone performs a “Negative” word, e.g. NOT, NONE, etc., without also performing a “Negative” nonmanual signal at the same time, this would also be ungrammatical. Thus, our system leverages these types of relationships between key classes of ASL manual signs and required ASL nonmanual signals to detect when the person in a video may be making an error. Since nonmanual signals often span multiple manual signs, and since movements of the head may require some onset/offset transitional time, many possible temporal alignments between the nonmanual and manual channels are possible, including various forms of overlap. Thus, the error-detection rules in our system use a temporal threshold to determine whether a grammatically-required nonmanual signal has occurred at the correct time, relative to the timing of the manual signs.

Table III lists the ASL grammatical errors recognized in our system. We categorize them into two groups: 1) Lexical errors happen with absence of proper facial expressions and head movements when certain classes of manual signs are performed. For instance, “NEG-Lexical” happens if a “Negative” manual sign is performed without a “Negative” facial expression or head movement. 2) Timing errors indicate that a grammatically important nonmanual signal happened but was too far from the clause boundary. The clause boundary is the temporal gap between two consecutive clauses. The students may submit a video of longer passages of ASL consisting of multiple sentences; in this case, a subset of inter-sign boundaries will be “clause boundaries.” Our system must therefore detect the clause boundaries throughout continuous videos. The timing errors capture how ASL signers must produce certain nonmanual signals in temporal alignment

TABLE II
CLASSES OF NONMANUAL SIGNALS OUR SYSTEM DETECTS.

Face Class	Head Pose (Head Class)	Eyebrows, Eyes, Nose
Negative	Shaking side-to-side (1)	Furrowed, Scrunched nose
YNQ	Tilted forward (2)	Raised
WHQ	Tilted forward (2)	Furrowed
RHQ	Tilted slightly to side (3)	Raised
Topic	Tilted slightly to side (3)	Raised, Widen eyes
Conditional	Tilted slightly to side (3)	Raised

TABLE III
ASL GRAMMAR ERRORS RECOGNIZED BY OUR SYSTEM.

Lexical errors: if required facial expressions or head movements are not found within S_L time from the sign. Timing errors (related to beginning and end): if facial expressions or head movements begin or end but with more than S_T time away from Clause Boundary.

Error Type	Manual Sign	Nonmanual Signal
WHQ-Lexical	WHQ word	Not (WHQ or RHQ)
YNQ-Lexical	YNQ word	Not YNQ
NEG-Lexical	Negative word	Not Negative
COND-Lexical	Conditional word	Not Conditional
YNQ-Beginning	Clause Boundary	YNQ begins
YNQ-End	Clause Boundary	YNQ ends
COND-Beginning	Clause Boundary	Conditional begins
TOPIC-Beginning	Clause Boundary	Topic begins

with the beginning or end of clauses. For instance, “YNQ-Beginning” happens if a “YNQ” nonmanual signal begins more than t seconds away from the closest clause boundary, where t is a given threshold.

C. System Overview and Contributions

As shown in Fig. 2, our framework includes 3D networks for multiple modalities: a *Hand Gesture Network*, a *Face Network*, and a *Head Network* to capture spatiotemporal information from hands, facial expressions, and head movements respectively. Then a sliding window-based approach in temporal dimension is applied to detect temporal boundaries of grammatically important gestures. A majority voting algorithm is designed to finalize the predictions for each frame by picking the most frequent detected class of gesture. A segment detection algorithm is implemented to connect the predictions belonging to the same class. The predictions are then pruned based on the confidence scores returned by the networks. Finally, our *Error Detection* algorithm detects grammatical mistakes by comparing the predictions from different modalities and analyzing their temporal correspondences.

The main contributions of the proposed framework are summarized as follows:

- To the best of our knowledge, this is the first framework for automatic detection of ASL grammatical mistakes in continuous signing videos.
- We propose a 3D multimodal framework to recognize grammatically important manual signs and nonmanual

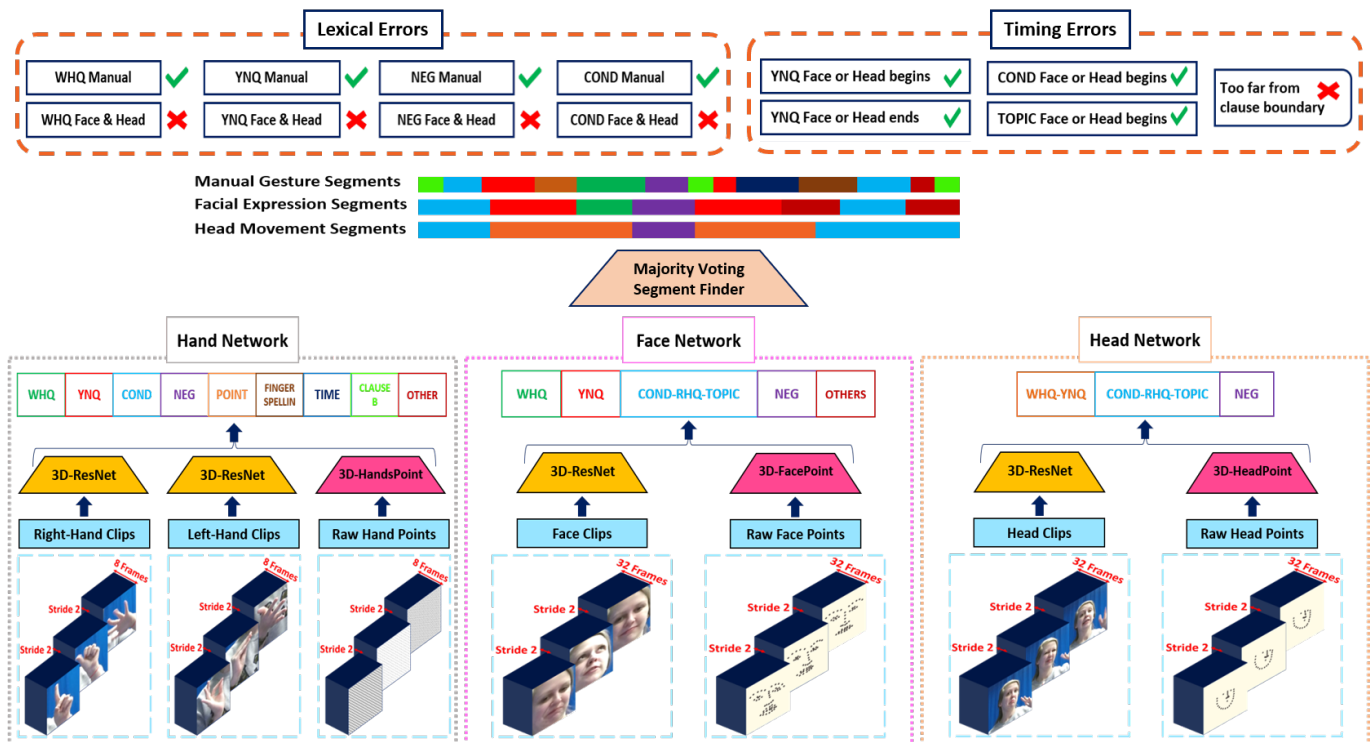


Fig. 2. The pipeline of the proposed multimodal 3DCNN framework for ASL nonmanual grammar recognition. The multiple modalities include hand gestures, facial expressions, and head movements. Raw coordinates and RGB clips of face and hands are fed to the networks to extract the pose and spatiotemporal information. Temporal boundaries of manual gestures and nonmanual signals are approximated using temporal sliding windows. The final predictions of frames are determined via Majority Voting algorithm and segments of all classes are extracted. Temporal correspondences of multiple modalities are compared to detect lexical and timing errors.

signals from continuous signing videos without temporal segmentation.

- We have collected a continuous sign language dataset named as ASL-HW-RGBD, consisting of 1,026 continuous videos of signing several sentences by fluent and student ASL signers. Our dataset covers different aspects of ASL grammar such as Conditional Sentences, Rhetorical Questions, WH Questions, YN Questions, Autobiography, Narrative, Pronouns, and Possessives. The dataset is annotated using ELAN annotation tool at the frame level.
- Our system is able to recognize grammatical elements from manual gestures, facial expressions, and head movements respectively. Furthermore, by analyzing the temporal correspondence of the grammar elements from multiple modalities, our system can effectively and efficiently detect 4 lexical errors as well as 4 timing errors (related to beginning and end).
- Our system generates instant feedback for ASL students about the detected grammatical mistakes in their continuous signing videos.

II. RELATED WORK

For continuous sign language recognition (CSLR), while most existing frameworks focus on the sentence-level translation of continuous signing, our system aims to recognize

ASL grammatical mistakes without necessarily recognizing each word in a sentence. The existing CSLR frameworks consist of three main stages: 1) temporal segmentation, 2) feature extraction from segments, and 3) sequence modeling to reconstruct the sentence. Our approach is mostly related to the first two stages as we do not need to fully translate ASL sentences to find grammatical elements.

Temporal Segmentation: Zhang *et al.* proposed a threshold matrix-based method combined with hidden Markov model (HMM) for coarse segmentation, followed by dynamic time warping (DTW) for fine segmentation [5]. Dawod *et al.* used contrast adjustment and motion detection analysis for segmentation process [6]. Yang *et al.* proposed dynamic time wrapping-based level building (LB-DTW) to segment sign sequences and then recognize signs [7]. Connectionist temporal classification (CTC) [8] is an end-to-end sequence learning model which is suitable for unsegmented input data to learn the correspondence between the input and output sequences and is used in [9], [10] for CSLR. Huang *et al.* proposed a novel method based on hierarchical attention network with latent space to eliminate the need for temporal segmentation [11]. Sliding window is a traditional approach generally used in dynamic event detection which is also employed for sign language recognition [12]. We employ a sliding window-based approach to detect temporal boundaries of manual gestures and nonmanual signals for grammar

analysis. Sliding windows cover all segments even subtle facial movements but the temporal boundaries are roughly approximated.

Video Feature Extraction: 3D convolutional neural networks (CNNs) are commonly used to capture spatiotemporal representations in videos. 3DCNN was first proposed for video action recognition in [13], and was improved in C3D [14] by using a similar architecture to VGG [15]. Recently, many 3DCNN models were proposed such as [16],[17],[18] and demonstrated effectiveness in video recognition tasks. One of the main challenges of training deeper networks is gradient vanishing and gradient explosion which was tackled in ResNet[19] by using *skip connections* and sending the previous feature map to the next convolutional block. 3D-ResNet[20] inherits the advantages of ResNet and also efficiently captures the temporal dimension. In this work, we employ 3D-ResNet [20] as the backbone network for video feature extraction due to its capability to capture spatiotemporal features in action recognition tasks.

III. APPROACH

We propose a multimodal framework to automatically recognize the grammatical errors from continuous signing videos without fully translating the entire sentences. Fig. 2 summarizes the pipeline of the proposed framework. To analyze the grammar in continuous signing videos, we first recognize grammatical elements from manual gestures, facial expressions, and head movements by employing 3D-CNN networks to capture spatiotemporal features. A sliding window-based approach in the temporal dimension is applied to detect temporal boundaries of grammatically important elements from all modalities. A majority-voting algorithm finalizes the prediction of each frame by considering the prediction of its neighboring frames. A segment-finder algorithm next connects the predictions belonging to the same class. Based on the frame-level predictions, we analyze the correspondence of grammatical elements in different modalities and detect two types of errors (Lexical and Timing) as well as correct signing. This section describes the data preprocessing, the design of hand-gesture, face, and head networks in detail, and our method for grammar-error recognition.

A. Data Preprocessing

We need to prepare the following components of the data for training: 1) Raw coordinates of face and hands key-points to extract pose information and dynamics of gestures. OpenPose [21] is employed to extract raw coordinates of face and hands from RGB videos. If hands or face are out of frame boundary, OpenPose returns zero values for the key-points. In such cases, we estimate the key-points coordinates in the missed frames with interpolation and extrapolation of previous and next frames. 2) Regions of Hands and head to extract the fine-grained spatiotemporal information. The extracted coordinates by OpenPose are used to crop head and hands regions from RGB images. 3) Registered face images to capture facial expressions without head movements. We first calculate a

mean face by averaging the coordinates of facial landmarks across the entire dataset, then for each frame, we use the facial landmarks (extract by OpenPose [21]) to register the face using affine transformation and with respect to the mean face. The raw face coordinates are also warped using the same process.

B. Network Architecture

Temporal Boundary Recognition: To recognize temporal boundaries of gestures in continuous videos, most existing methods used temporal segmentation of signs as a preprocessing step. However, incorrect temporal segmentation can consequently lead to wrong recognition of manual gestures and nonmanual signals thus negatively affect the entire pipeline by generating too many false positives or false negatives. Instead of using temporal segmentation, we add an extra “*Others*” class to the manual and facial gesture classes, which includes all instances that do not belong to any of the grammatically important categories. We then use a sliding-window approach to detect the intervals with high confidence scores of grammar elements. It is worth noting that this part is only used during testing where no temporal boundaries are provided. In the training phase, the annotations of temporal boundaries are used.

Backbone network for video feature extraction: We design three networks (Hand Gesture, Face, and Head) to extract video features. All of them use same backbone 3D-ResNet (with 34 layers) architecture with 5 convolutional blocks. The first one consists of one convolutional layer, followed by batch normalization, ReLU, and a max-pooling layer. The other 4 convolutional blocks are 3D residual blocks with skip connections. The number of kernels in the five convolutional blocks are {64, 64, 128, 256, 512}. The Global Average Pooling (GAP) is followed after the 5th convolutional block to produce a 512-dimensional feature vector. Then, one fully connected layer is applied to produce the final prediction. All the networks are optimized with Weighted Cross Entropy loss and Stochastic Gradient Descent (SGD) optimizer.

C. Hand Gesture Network

As described in Table I, the hand gesture network is trained to classify manual signs into 9 classes related to our system’s error-detection rules: *Conditional*, *Yes-No-Question (YNQ)*, *WH-Question (WHQ)*, *Negative*, *Time*, *Pointing*, *clause boundary*, *Finger-Spelling*, and *Others*. The class “*Others*” includes all ASL signs that do not belong to the first 8 classes. The gesture network consists of three streams: RGB clip of right hand, RGB clip of left hand, and raw coordinates of hands. All streams are jointly trained in an end-to-end manner. The feature vectors from three streams are concatenated and fed to Softmax for classification.

3D-ResNets for Hands extract spatiotemporal information of right hand and left hand clips by employing 3D-ResNet architecture with 34 layers and temporal duration 8, each 3D-ResNet producing a 512-dimensional feature vector.

3D-HandsPointNet is designed to extract pose information and dynamics of hand gestures from raw coordinates of hands

keypoints. It consists of two fully connected layers, each one followed by batch normalization and ReLU. The input is a $t \times 90$ matrix where t is the temporal duration (set to 8 frames in our system). Each column of the matrix is corresponding to one frame of the video and is a 90-dimensional vector consisting of x, y coordinates of 21 key points of right hand, 21 key points of left hand as well as the centers of right hand, left hand and face (to capture the relative positions). The coordinates are normalized using mean and standard deviation from the dataset. The dimension of the first and second fully connected layers are 90×40 and 40×20 , respectively. The columns of the output matrix ($t \times 20$) are concatenated resulting in a 160-dimensional feature vector.

D. Face Network

As described in Table II, the face network is designed to classify facial expressions to 5 classes: *Conditional-Topic-RHQ*, *Negative*, *Yes-No-Question (YNQ)*, *WH-Question (WHQ)*, and *Others*. The *Conditional-Topic-RHQ* class combines Conditional, Topic, and RHQ facial expressions as one class for classification due to their high similarity. The class *Others* includes all facial expressions that are not among the first 4 classes (not related to our system’s error-detection rules). Face network consists of two streams: whole face region and coordinates of facial key points which are jointly trained and their feature vectors are concatenated.

3D-ResNet for Facial Expressions receives RGB clips of registered face images as input resulting in a 512-dimensional feature vector. The purpose of face registration is to remove the head movements so that the network captures only facial expression changes. The temporal duration of this network is 32 which is longer than Hand Gesture network because facial expressions typically last longer.

3D-FacePointNet takes clips of registered raw face coordinates as input, i.e., a $t \times 106$ matrix where t is the temporal duration (set to 32 in our system). Each column corresponds to one frame and is 106-dimensional vector including x, y coordinates of 53 facial key points extracted by OpenPose [21] (all the facial key points such as eyes, eyebrows, nose, mouth, except face contour). The network consists of two fully connected layers, each one followed by a batch normalization and a ReLU. The input matrix is spatially compressed to $t \times 20$ by the first fully connected layer. It is then transposed ($20 \times t$) and temporally compressed to 20×10 by the second fully connected. The columns are then concatenated resulting in a 200-dimensional feature vector.

E. Head Network

The head network extracts spatiotemporal information of head movements. As listed in Table II, the head movements are categorized to three classes: *Shaking head side-to-side (Negative)*, *Head tilted forward (YNQ, WHQ)*, and *Head tilted slightly to side (Topic, Conditional, RHQ)*. The head network also has two streams and is similar to face network with subtle differences.

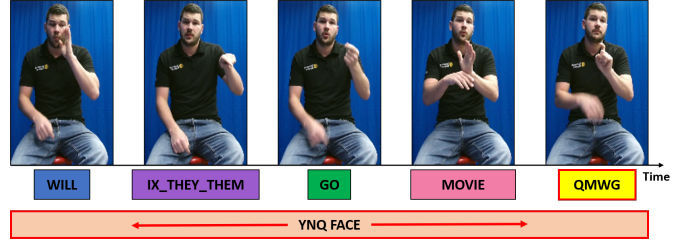


Fig. 3. An example of temporal alignment of signs in a Yes-No Question (YNQ) in ASL for the sentence “Will they go to a movie?”. The YNQ facial expression happens during the entire question while the YNQ-related manual sign (QMWG) happens only at the end.

The first stream is 3D-ResNet receiving clips of head images. The head regions are wider than face regions and are not registered. The second stream is 3D-HeadPointNet, receiving clips of raw coordinates of head as input which are 58-dimensional vectors including 29 key points of face (face contour, nose, center of eyes and center of mouth) for each frame. 3D-HeadPointNet consists of two fully connected layers to reduce temporal and spatial dimensions and is similar to 3D-FacePointNet network. Likewise, the input from two streams are concatenated and fed to Softmax.

F. Grammatical Error Recognition

Our method for recognizing ASL grammar errors contains two main steps: ASL segment detection and grammatical error recognition. Note that the grammatical error recognition is only conducted in testing phase.

ASL Segment Detection: In order to detect the temporal boundaries of gestures, a sliding window is applied throughout the input video, and the frames of each window are fed to the networks. In our experiments, the sliding-window size is 8 for the hand-gesture network, and 32 for the face and head networks, with a stride at 2. The lengths of sliding windows are set based on the average duration of manual or facial gestures in the proposed dataset. Denoting the size of sliding window by S , there are $\lceil \frac{n-S+1}{2} \rceil$ sliding windows for a video with n frames. To speed-up the process, at each step, a batch of 64 sliding windows are fed to the network in parallel for testing.

Each video frame belongs to several overlapping sliding windows resulting in a list of predictions for each frame. A Majority Voting method is implemented to determine the final prediction for each frame, i.e. by selecting the most frequent predicted labels in the list and referring to prediction of previous frames if there is a tie. The pseudocode is illustrated in *Algorithm 1*.

After obtaining the final prediction for each frame via Majority Voting algorithm, consecutive frames with the same label (i.e. segments) are detected from manual gestures, facial expressions, and head movements independently. For the frames in each segment with the same prediction, if confidence score of at least one frame is greater than the specified threshold, the prediction is kept without change. Otherwise, the predicted label is changed to class *Others*. In our experiments, the confidence score threshold is set to 0.8 for all modalities.

Algorithm 1 Majority Voting Algorithm

```
1: procedure
2:    $i \leftarrow$  current frame index
3:    $P \leftarrow$  final predictions for previous frames
4:    $l \leftarrow$  list of predictions for frame  $i$ 
5:    $T \leftarrow$  empty hash table
6:   loop over  $l$ :
7:     if class  $C \in T.keys$  then  $T[C] += 1$ 
8:     else  $T[C] = 1$ 
9:      $\triangleright T$  keeps frequency of predicting  $C$  as label of  $i$ .
10:  candidates:
11:    $K = [ C \in T.keys \text{ if } T[C] = \max(T.values) ]$ 
12:    $\triangleright K$  is a list of predictions with maximum frequency.
13:   if  $|K| = 1$  then return  $K[0]$ 
14:    $j \leftarrow 1$ 
15:   while  $j \leq 3$  do
16:     if  $P[i - j] \in K$  then return  $P[i - j]$ 
17:      $j \leftarrow j + 1$ 
18:    $\triangleright$  Return prediction of previous frames if belong to  $K$ .
19:   return a random element of  $K$ 
```

ASL Grammatical Error Recognition: ASL grammar errors are recognized based on a set of rules related to the synchronization between the manual gestures, facial expressions, and head movements, as listed in Table III. The facial expressions (and head movements) in ASL are usually longer than the manual gestures, making the temporal alignments more complicated. Fig. 3 shows an example of temporal alignment for a Yes-No Question (YNQ) while the signer is performing a sentence consisting of five manual signs: “WILL,” “IX-THEY-THEM,” “GO,” “MOVIE,” and “QMWG” (question mark wiggling sign). Rather than performing the YNQ nonmanual signal only during QMWG, the signer performs it during the entire sentence (which is also grammatically acceptable). Our system looks for potential lexical and timing errors (see definitions in Table III), and sets temporal proximity thresholds of 200 msec for lexical errors and 1 second for timing errors.

IV. CONTINUOUS SIGN LANGUAGE DATASET

We have collected a sign language dataset, ASL-HW-RGBD, consisting of 1,026 continuous videos (some consisting of multiple ASL sentences), as produced by 46 ASL signers, including both fluent and student signers. Each video was produced as a response to a homework-style assignment. Some assignments simply ask students to translate a set of English sentences into ASL, and others ask students to invent short multi-sentence passages. Each assignment is designed to elicit sentences that consist of several ASL grammar elements: For instance, the first homework is mainly about signing two basic question types in ASL, which are WH-Questions (what, who, where, how, etc) and Yes-No-Questions. The second homework asks students to compose longer multi-sentence videos with autobiographical details.

The third homework encourages students to use various pointing signs and fingerspelling, while the fourth homework encourages students to create a longer video with multiple items under discussion and questions asked to the camera. The fifth homework is centered around pronouns, negation, and questions. Finally, the sixth homework is about conditional sentences and rhetorical questions.

The dataset has been annotated using the ELAN annotation tool, with several parallel timeline tiers of annotations for each video (in which spans of time in the video are labeled with key linguistic information). We briefly list some of the most important tiers for our research: 1) Clause: beginning and end of the clause, 2) Fingerspelling: spelling out words by using hand shapes that correspond to the letters of the word, 3) Lexical Pointing: referring to person or concept under discussion, 4) Wanted Words: words that may be used more frequently than others including several classes such as WH-Question, Yes-No Question, Negative, Conditional, Time, and Pointing, 5) Facial Expressions including several classes such as Conditional, Negative, Yes-No Question, WH-Question, Rhetorical Question, and Topic.

V. EXPERIMENTS

A. Implementation Details

Our proposed models are implemented in PyTorch on four Titan X GPUs. To avoid over-fitting, our models are pretrained on Kinetics [22] which is a large human action dataset. The original resolution of RGB videos in our ASL-HW-RGBD dataset is $1,920 \times 1,080$ pixels. To prepare the input for hand gesture and head networks, we use the coordinates extracted by OpenPose [21] to crop 240×240 regions around the center of face and 160×160 regions around the center of each hand. The 240×240 face images are used for face registration, and are further cropped to 100×128 for face network. All these bounding boxes are adjusted based on the mean and variance of key-points coordinates in the dataset to guarantee the face regions are inside the 100×128 bounding box. We finally resize all input images to 134×134 for training and testing. Data augmentation techniques such as random cropping and random rotation are used during training. In every iteration of training, 112×112 image patches are randomly cropped from 134×134 input images for data augmentation. Random rotation (with a degree randomly selected in a range of $[-10, 10]$) is applied on the cropped patches to further augment the data. During the testing, only the center patches of size 112×112 are used for predictions. The models are fine-tuned for 200 epochs with an initial learning rate of $\lambda = 3 \times 10^{-3}$, reduced by a factor of 10 every 50 epochs. The batch size is 128.

B. Grammar Element Detection Results

The grammar-error recognition relies on the detection accuracy of the manual signs and nonmanual signals. A wrong prediction of a manual gesture would consequently leads to search for a non-essential element in other modalities, resulting in a false positive grammar error. Moreover, bias towards predicting a certain class of face or head can lead

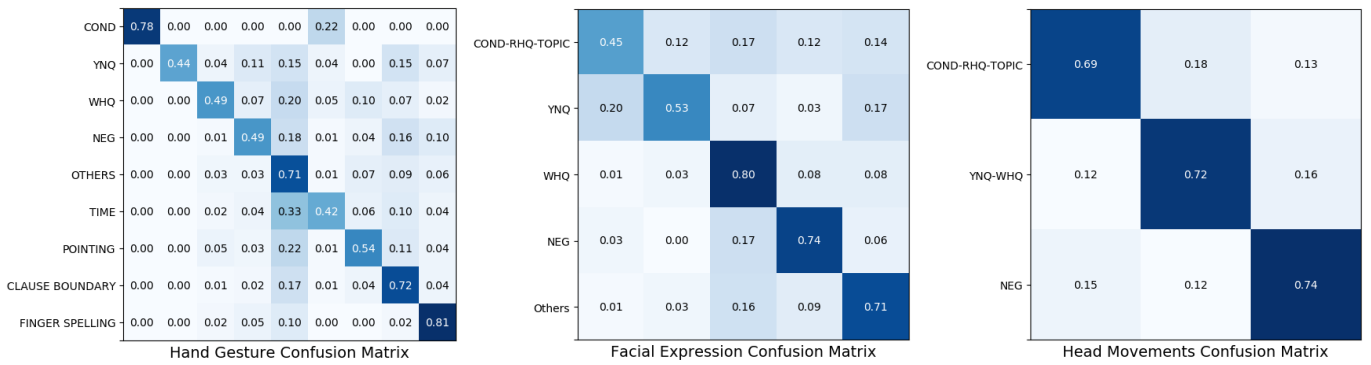


Fig. 4. Normalized confusion matrices of Hand Gesture Network with 66.84% accuracy (left), Face Network with 66.32% accuracy (middle), and Head Network with 71.14% accuracy (right). The rows are ground truth labels and the columns are predictions.

to missing grammatical mistakes when that class is required but has not been performed, resulting in a false negative. For example, suppose the face network were to predict the facial expression as “Negative” in most cases. When the signer is performing a “Negative” manual sign (hand-gesture) but omitted the “Negative” face, the system may wrongfully predict the facial expression as “Negative.” Therefore, the signer’s “NEG-Lexical” error (see Table III) would not be detected.

Normalized confusion matrices of Hand-Gesture, Face, and Head Networks are shown in Fig. 4. The value in cell (i, j) of each matrix is the percentage of predicting ground truth label i as class j . We observe that misclassifications are mostly caused by class *Others*. The reason is that class *Others* is heterogeneous including all instances that do not belong to rest of categories, but may contain similar manual gestures and nonmanual signals. Therefore, it is easy to be confused with the rest of classes.

C. Grammatical Error Recognition Results

To evaluate our framework, we test a subset of videos in our ASL-HW-RGBD dataset with labeled grammatical errors listed as Table III. The grammatical errors in these clips are detected and annotated by ASL linguists beforehand. Note that in the training phase, the input data does not include grammatical errors at all. For testing, we first recognize the grammatical elements from multiple modalities and then analyze grammar errors independently. A lexical error is detected whenever a grammatically important hand gesture is recognized (with confidence above C) but the corresponding facial expression and head movement are not detected within S_L msec of the hand gesture. Timing errors are detected if there is a grammatically important facial expression or head movement (with confidence score above C) but it is performed more than S_T seconds from the closest clause boundary. To reduce false-positive and false-negative errors, and based on the performance of our networks, we set the confidence score C to 0.8, S_L to 200 msec, and S_T to 1 second, to achieve reliable detection. The results of grammatical error recognition are shown in Table IV. The second column includes the ground truth number of instances and the third column is the number

of instances recognized by our system. The last column (True Positive Rate) is the ratio of the recognized errors to ground truth errors. For *Error-COND-Beginning*, *Error-YNQ-Beginning* and *Error-YNQ-End*, there are not enough ground truth instances so they are not included in the evaluation. The average true positive rate is 60%. The true positive rate of *Error-COND-Lexical* is the highest among lexical errors because “Conditional” hand gestures can be recognized with relatively high accuracy (78%) compared to other classes. On the other hand, the recognition accuracy of “YNQ”, “WHQ”, and “Negative” hand gestures are less than 50% (see Fig. 4) and can be missed in some cases resulting in false negative error detection. Furthermore, the true positive rate of *Error-TOPIC-Beginning* is the highest because recognition of the “Topic” class either from face or head leads to exploring the possibility of this error and with relatively high recognition accuracy of “Clause Boundary” from hand gestures, this error can be correctly recognized in most cases. Most false-positive errors related to “Error-NEG-Lexical.” The reason is that 11% of YNQ manual gestures, and 7% “WHQ” manual gestures, were wrongfully recognized as being in the “Negative” class, as shown in the hand-gesture confusion matrix in Fig. 4. In these cases, there were no true instances of “Negative” gesture, therefore, the system did not predict a “Negative” face or head. As a result, a false positive “Error-NEG-Lexical” had been detected. In future work, we aim to reduce false-positive errors by improving the recognition results of grammatical elements in various modalities.

TABLE IV
ASL NONMANUAL GRAMMATICAL ERROR RECOGNITION RESULTS.

Error Type	Ground Truth	Recognized	TP Rate
Error-YNQ-Lexical	12	7	58.3%
Error-NEG-Lexical	36	14	38.8%
Error-WHQ-Lexical	40	23	57.5%
Error-COND-Lexical	24	18	75.0%
Error-TOPIC-Beginning	23	19	82.6%
Total	135	81	60.0%

Our system is able to generate feedback for students on average in less than 2 minutes for 1 minute ASL videos

including all the steps from pre-processing the video, testing the networks, and error detection. Our system has also been tested on videos recorded with cellphones, and preliminary analysis suggests it is robust to varying conditions, e.g. illumination, image quality, and distance from the camera.

VI. CONCLUSION

In this paper, we have proposed a 3DCNN-based multimodal framework to automatically recognize ASL grammatical elements from continuous signing videos. Then a sliding window-based approach is applied to detect 8 types of ASL grammatical mistakes by checking the temporal correspondence between manual signs and nonmanual signals in signing videos. Our system generates instant feedback for ASL students about the grammatical aspects of their performance without fully translating the sentences. We have collected and annotated a new dataset, *ASL-HW-RGBD*, consisting of 1,026 continuous sign language videos. Our system is able to recognize 60% of grammatical mistakes. Our future work will aim to develop more advanced methods to handle the complex relations between manual gestures and nonmanual signals in continuous ASL videos to improve the recognition accuracy of ASL grammatical elements as well as ASL grammar errors.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under award numbers 1400802, 1400810, and 1462280.

REFERENCES

- [1] N. Furman, D. Goldberg, and N. Lusin, "Enrollments in languages other than english in united states institutions of higher education, fall 2010," *Retrieved from http://www.mla.org/2009_enrollmentsurvey*, 2010.
- [2] L. Jing, E. Vahdani, M. Huenerfauth, and Y. Tian, "Recognizing american sign language manual signs from rgb-d videos," *arXiv preprint arXiv:1906.02851*, 2019.
- [3] R. B. Wilbur, "Effects of varying rate of signing on asl manual signs and nonmanual markers," *Language and speech*, vol. 52, no. 2-3, pp. 245–285, 2009.
- [4] D. M. Perlmutter, "What is sign language?" *Linguistic Society of America*, vol. 1325, 2011.
- [5] J. Zhang, W. Zhou, and H. Li, "A threshold-based hmm-dtw approach for continuous sign language recognition," in *Proceedings of International Conference on Internet Multimedia Computing and Service*, 2014, pp. 237–240.
- [6] A. Y. Dawod, M. J. Nordin, and J. Abdullah, "Gesture segmentation: automatic continuous sign language technique based on adaptive contrast stretching approach," *Middle-East Journal of Scientific Research*, vol. 24, no. 2, pp. 347–352, 2016.
- [7] W. Yang, J. Tao, and Z. Ye, "Continuous sign language recognition using level building based on fast hidden markov model," *Pattern Recognition Letters*, vol. 78, pp. 28–35, 2016.
- [8] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [9] R. Cui, H. Liu, and C. Zhang, "Recurrent convolutional neural networks for continuous sign language recognition by staged optimization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7361–7369.
- [10] J. Pu, W. Zhou, and H. Li, "Iterative alignment network for continuous sign language recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4165–4174.
- [11] J. Huang, W. Zhou, Q. Zhang, H. Li, and W. Li, "Video-based sign language recognition without temporal segmentation," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [12] E.-J. Ong, O. Koller, N. Pugeault, and R. Bowden, "Sign spotting using hierarchical sequential patterns with temporal intervals," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1923–1930.
- [13] S. Ji, W. Xu, M. Yang, and K. Yu, "3d convolutional neural networks for human action recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 1, pp. 221–231, 2013.
- [14] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 4489–4497.
- [15] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [16] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*. IEEE, 2017, pp. 4724–4733.
- [17] A. Diba, M. Fayyaz, V. Sharma, A. H. Karami, M. Mahdi Arzani, R. Yousefzadeh, and L. Van Gool, "Temporal 3D ConvNets: New Architecture and Transfer Learning for Video Classification," *ArXiv e-prints*, Nov. 2017.
- [18] Z. Qiu, T. Yao, and T. Mei, "Learning spatio-temporal representation with pseudo-3d residual networks," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [19] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [20] K. Hara, H. Kataoka, and Y. Satoh, "Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet?" in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 6546–6555.
- [21] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh, "OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields," in *arXiv preprint arXiv:1812.08008*, 2018.
- [22] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev *et al.*, "The kinetics human action video dataset," *arXiv preprint arXiv:1705.06950*, 2017.