



Special Section on 3DOR 2021

SHREC 2021: Skeleton-based hand gesture recognition in the wild

Ariel Caputo^{a,*}, Andrea Giachetti^a, Simone Soso^a, Deborah Pintani^a, Andrea D'Eusanio^b, Stefano Pini^b, Guido Borghi^c, Alessandro Simoni^b, Roberto Vezzani^b, Rita Cucchiara^b, Andrea Ranieri^d, Franca Giannini^d, Katia Lupinetti^d, Marina Monti^d, Mehran Maghoumi^e, Joseph J. LaViola Jr^e, Minh-Quan Le^f, Hai-Dang Nguyen^f, Minh-Triet Tran^f

^a Department of Computer Science, University of Verona, Italy^b Università di Modena e Reggio Emilia, Dipartimento di Ingegneria "Enzo Ferrari"^c Università di Bologna, Dipartimento di Informatica - Scienza e Ingegneria^d Istituto di Matematica Applicata e Tecnologie Informatiche "Enrico Magenes"^e University of Central Florida, Department of Electrical Engineering and Computer Science^f University of Science, Ho Chi Minh City, Vietnam

ARTICLE INFO

Article history:

Received 1 April 2021

Revised 5 July 2021

Accepted 6 July 2021

Available online 10 July 2021

Keywords:

Gesture recognition

Hand skeleton

Online

Interaction

ABSTRACT

Gesture recognition is a fundamental tool to enable novel interaction paradigms in a variety of application scenarios like Mixed Reality environments, touchless public kiosks, entertainment systems, and more. Recognition of hand gestures can be nowadays performed directly from the stream of hand skeletons estimated by software provided by low-cost trackers (Ultraleap) and MR headsets (Hololens, Oculus Quest) or by video processing software modules (e.g. Google Mediapipe). Despite the recent advancements in gesture and action recognition from skeletons, it is unclear how well the current state-of-the-art techniques can perform in a real-world scenario for the recognition of a wide set of heterogeneous gestures, as many benchmarks do not test online recognition and use limited dictionaries. This motivated the proposal of the SHREC 2021: Track on Skeleton-based Hand Gesture Recognition in the Wild. For this contest, we created a novel dataset with heterogeneous gestures featuring different types and duration. These gestures have to be found inside sequences in an online recognition scenario. This paper presents the result of the contest, showing the performances of the techniques proposed by four research groups on the challenging task compared with a simple baseline method.

© 2021 Elsevier Ltd. All rights reserved.

1. Introduction

The recognition of gestures based on hand skeleton tracking is becoming the default interaction method for the new generation of Virtual Reality (VR) and Mixed Reality (MR) devices like Oculus Quest and Microsoft Hololens, implementing specific, advanced solutions [5,16]. Low-cost hand tracking devices with good performances are available since 2010 [19] and are used in several ap-

plication domains and research works. Real-time hand pose tracking is now possible from single-camera input using Google tools [22]. It is, therefore, extremely likely that most of the future hand gesture recognition tools will work directly on the hand skeleton poses and not on RGB or depth images. These facts strongly motivate research efforts aimed at the development of such tools. In practical application scenarios, these gesture recognizers need to work in real-time and to be able to detect and correctly label gestures "in the wild" within a continuous sequence of hand movements.

Several methods have been recently proposed in the literature for the skeleton-based gesture recognition task. However, as pointed out in [1], current available benchmarks that focus on online-recognition scenarios are limited. Many of them do not test recognizers in an online setting or evaluate the methods on limited vocabularies and not including many gesture types. Hand gestures, in fact, can be classified into different types according to their distinctive features. Some gestures are static, characterized by keeping a fixed hand pose for a minimum amount of time. Others are dy-

* Corresponding author.

E-mail addresses: ariel.caputo@univr.it (A. Caputo), andrea.giachetti@univr.it (A. Giachetti), simone.soso@studenti.univr.it (S. Soso), deborah.pintani@studenti.univr.it (D. Pintani), andrea.deusanio@unimore.it (A. D'Eusanio), spini@unimore.it (S. Pini), guido.borghi@unibo.it (G. Borghi), alessandro.simoni@unimore.it (A. Simoni), roberto.vezzani@unimore.it (R. Vezzani), rita.cucchiara@unimore.it (R. Cucchiara), andrea.ranieri@cnr.it (A. Ranieri), franca.giannini@cnr.it (F. Giannini), katia.lupinetti@cnr.it (K. Lupinetti), marina.monti@cnr.it (M. Monti), maghoumi@knights.ucf.edu (M. Maghoumi), jjl@eeecs.ucf.edu (J.J. LaViola Jr), lmquan@selab.hcmus.edu.vn (M.-Q. Le), nhdang@selab.hcmus.edu.vn (H.-D. Nguyen), tmtriet@fit.hcmus.edu.vn (M.-T. Tran).

<https://doi.org/10.1016/j.cag.2021.07.007>

0097-8493/© 2021 Elsevier Ltd. All rights reserved.

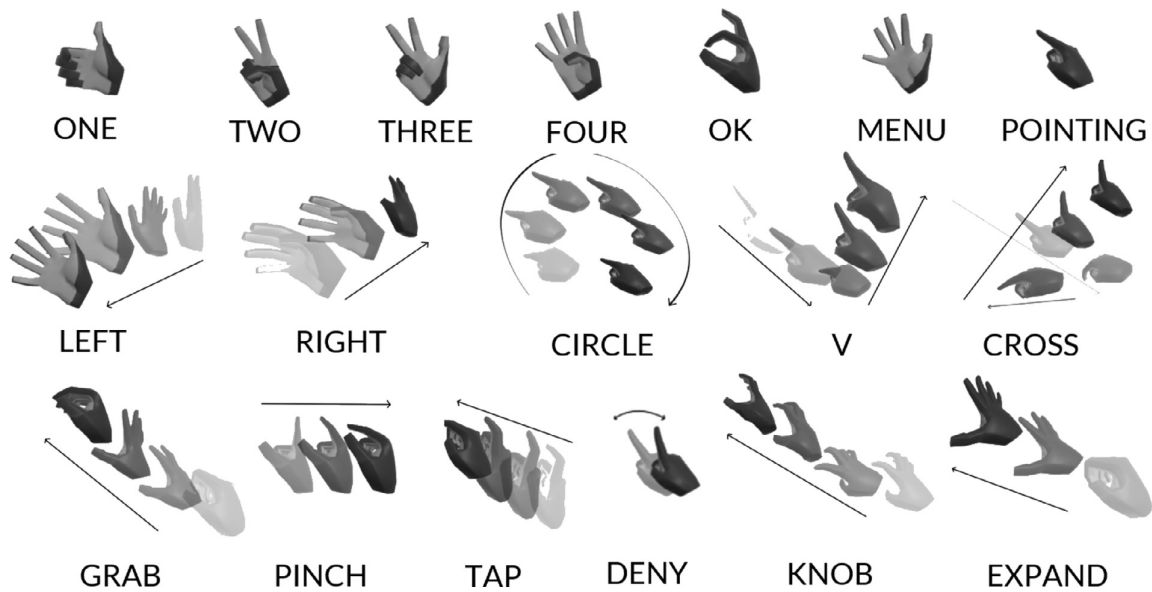


Fig. 1. Gesture templates for the different classes. Top row: static gestures. Middle row: coarse dynamic gestures. Bottom row: fine dynamic gestures.

dynamic and characterized by a single trajectory with the hand pose that does not change or it is not semantically relevant. Others are dynamic and characterized not only by a global motion, but also by the evolution of fingers' articulation over time.

Previous contests organized on skeleton-based gesture recognition were limited to offline recognition (SHREC'17 Track: 3D Hand Gesture Recognition Using a Depth and Skeletal Dataset [3]) or featured a very limited dictionary of gestures (SHREC 2019 track on online gesture detection [2]).

For this reason, we created a novel dataset including 18 gesture classes belonging to different types. A subset of 7 classes are static, characterized by a hand pose kept fixed for at least one second (One, Two, Three, Four, OK, Menu, Pointing). The remaining ones are dynamic, 5 coarse, characterized by a single global trajectory of the hand (Left, Right, Circle, V, Cross) and 6 fine, characterized by variations in the fingers' articulation (Grab, Pinch, Tap, Deny, Knob, Expand). Fig. 1 shows the gestures' templates. A peculiarity of the data collected is that the gestures are executed within long sequences of hand gesticulation, as they were captured during generic user interaction.

Given the dataset, we proposed an online recognition task within the Eurographics SHREC 2021 framework. This paper reports on the outcomes of the contest's result. The paper is organized as follows: Section 2 presents the novel dataset, Section 3 the proposed task and the evaluation method, Section 4 presents the groups participating in the contest and the methods proposed together with a baseline method.

2. Dataset creation

The dataset created for the contest is a collection of 180 gesture sequences. Each sequence, captured using a Leap Motion Device, features either 3, 4, or 5 gestures interleaved with non-significant gesticulation. The dictionary used consists of 18 classes of gestures, each appearing in the dataset an equal amount of time (40 occurrences per class).

Gestures were performed by five different subjects in pre-determined sequences. The execution of the dictionary gestures followed specific templates shown in advance. Non-significant ges-

tication was limited to a restricted set of allowed movements to avoid biases on false detections.

We designed a randomized set of 36 sequences (12 with 3 gestures, 12 with 4 gestures, and 12 with 5 gestures). Each set thus includes 8 samples of the 18 gesture classes (a total of 144 gesture samples across the 36 sequences). Each subject recorded such a set of sequences for a total of 180 sequences with 720 gestures.

Three of the acquired sets (108 sequences with 24 samples of each gesture class) were given to the participants as the training set, with the associated annotations of gestures begin/end and labels. The remaining two (72 sequences with 16 occurrences of each gesture class) were given to the participants as the test set without associated labels. Gestures in the test sets were performed by subjects not involved in the training set creation.

We performed the acquisition of the sequences using a simple Unity app running on a desktop PC. Subjects had to wear a head-band with the leap motion device mounted over simulating sensor mounted on VR/AR glasses. The graphical interface of the app suggested the randomized temporal sequences of the gestures.

The accurate annotation of the timestamps of the start and end of the gestures was subsequently manually performed with another specific Unity application. In particular, the annotation tool allows to navigate the sequences, showing frame-by-frame the hand's movement. The application then allows to mark frames as the beginning or end of a gesture.

3. Task proposed and evaluation

The goal of the participants was to detect correctly the gestures included in the sequence with an online detection approach. The gesture database captured was split as described into a training set with associated annotations of gesture time stamp and labels, that could be used to train the detection algorithms and a test set with no annotations available. Participants had to provide a list of the gestures detected in the test set with associated labels, start time stamps and end timestamps.

The results have been evaluated using different metrics. First, we use the Jaccard Index as proposed in other online gesture recognition contests [18,23] to measure the average relative overlap between the ground truth and the predicted label sequences

for the gesture sequences recording. If $GT_{s,i}$ is a vector corresponding to the time sampling of the gesture with label i filled with ones where the gesture is actually performed and zeros elsewhere, and $P_{s,i}$ is the corresponding prediction, the index is given by:

$$J_{s,i} = \frac{GT_{s,i} \cap P_{s,i}}{GT_{s,i} \cup P_{s,i}} \quad (1)$$

Following our previous SHREC contest on Online Gesture recognition [2], we defined specific metrics to obtain an estimation of how the specific methods can measure the performances of a recognizer in a typical use context, where we are interested in detecting correctly gestures with a short delay after their execution and avoid false detections. We, therefore, counted the "detection rate" in the test data, e.g. the percentage of predicted gestures (of each class) corresponding to ground truth ones correctly detected. A prediction is considered corresponding to the ground truth if it has a temporal intersection ratio with the ground truth one higher than 0.5 and the same class label. We then measure also the false-positive ratio, i.e. the ratio between the number of gestures (of a particular class) predicted and not corresponding to ground truth ones divided by the total number of gestures of that class in the sequences.

4. Participants and methods

Five research groups were registered for the contest and sent results, but one retired after the evaluation. Each group sent up to three annotation files that were obtained with different methods or parameters' settings. The methods are described in the following subsection, together with the simple technique that we used as baseline.

4.1. Baseline: Dissimilarity-based classification

As a baseline method, we customized an algorithm used in [1] based on class-specific binary classifiers trained with dissimilarity features and a sliding window approach for online detection. We created a gesture dictionary with the segmented labeled intervals cropped from the training sequence and a set of "non-gesture" examples randomly cropped from non-labeled parts of the sequences. Each gesture has been re-sampled to 20 time steps. Half of the data have been used as the representation set for the dissimilarity vectors estimation.

Given the re-sampled gestures, we define four sets of dissimilarity vectors:

- palm trajectory dissimilarity: for each gesture of the representation set, we estimate the sum of the Euclidean distances of the corresponding points in query one. Given N gestures in the representation set, we get N features for the query gesture descriptor.
- hand articulation dissimilarity: first we estimate the evolution of distances between adjacent fingertips and between fingertips and the palm keypoint (Fig 2). We then calculate 9 dissimilarity components as the sums over corresponding time samples of the differences between the values of the 9 distances in the query and in the representation set gestures.
- palm trajectory length dissimilarity: the difference in length between the query gesture and the representation set gestures.
- palm velocity dissimilarity: the sum of the differences of the velocity magnitude samples at corresponding time steps,

For each gesture class, we train class-specific linear SVM classifiers using binary labels (gesture vs non-gesture).

During the online gesture recognition, we use a sliding window approach. To detect gestures at time t we give as input for each trained classifier, the hand pose samples cropped in the window

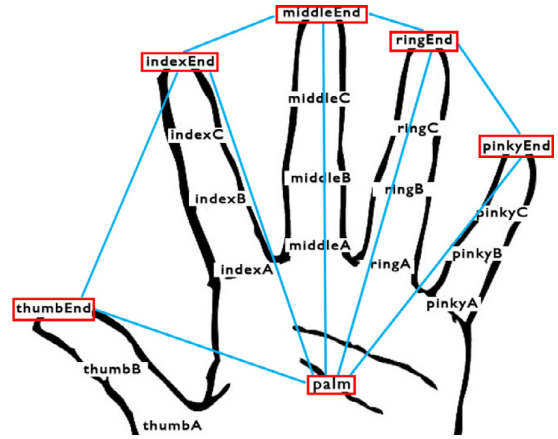


Fig. 2. Keypoints of the hand's skeleton (red) used to calculate the features for the dissimilarity-based method (4.1). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

covering the interval $[t - l(c), t]$, where $l(c)$ is the average duration, in frames, of the gestures of class c in the training set, re-sampled in 20 steps. Windows are sampled every 6 frames of the original sequence. If a gesture is detected, we assign it a duration equal to $l(c)$. If a gesture is detected in multiple consecutive frames the predictions are merged and the estimated duration is incremented by the number of consecutive detections multiplied by the windows sampling step.

4.2. Group 1: Transformer network based method

4.2.1. Method description

Group 1 proposed a dynamic gesture recognition system based on the Transformer model [17]. The framework is based on different type and combinations of features. The first group of features are provided by the *Leap Motion SDK* and consist in the 3D position of the hand joints. They enrich this features computing the joint velocity and acceleration. At time t , given the sequence of the 3D position of the i -th joint $J_i^t = (x_i^t, y_i^t, z_i^t)$, speed \mathbf{s} and acceleration \mathbf{a} are computed following these formulas:

$$\mathbf{s}_i^t = [x_i^t - x_i^{(t-1)}, y_i^t - y_i^{(t-1)}, z_i^t - z_i^{(t-1)}]$$

$$\mathbf{a}_i^t = [x_i^t - 2x_i^{(t-1)} + x_i^{(t-2)}, y_i^t - 2y_i^{(t-1)} + y_i^{(t-2)}, z_i^t - 2z_i^{(t-1)} + z_i^{(t-2)}]$$

Moreover, at time t the joint-to-joint 3D distances expressed as a matrix D of size $3 \times N \times N$, where N is the total number of the hand joints, is also computed. Each element $d_{k,j,z} \in D$ is computed as:

$$d_{k,j,z} = \sqrt{(J_{k,z}^t - J_{j,z}^t)^2}, \quad k, j \in N, \quad z \in [0, 3] \quad (2)$$

Therefore, the final dimension of the input feature vector depends on the types of the feature used. A feature vector composed by the hand features provided by the *Leap Motion* device and the speed and acceleration has size of 240. Including the 3D position and rotation of each hand joint the feature vector has a size of 640.

As pre-processing, Group 1 tested two different normalization techniques. In the first one, they normalized the joint positions in a per-instance manner to obtain a per-axis zero mean and unit variance. In the second one, they additionally divided each joint position by the hand "size", set as the distance between the joints *indexA* and *pinkyA*, before the zero mean and unit variance normalization.

At inference time, the network predicts the current gesture on a single-step time scale and a *Finite State Machine* (FSM) is used to

detect the beginning and the end of each gesture. Group 1 provided three different version of the presented algorithm: in the first and second cases, the feature vector is composed of the hand feature and speed and acceleration and then normalized using the two operations above mentioned. In the third case, all types of features normalized with zero mean and unit variance are included.

4.2.2. Model architecture

The proposed method is composed of a transformer module [17] (with internal dropout), followed by a fully connected layer applied to each step output which predicts the gesture class (including the class “no-gesture”).

Formally, the model can be defined as:

$$Y(\mathbf{x}) = F(\text{Encoders}(\mathbf{x} + PE)) \quad (3)$$

where $F(\cdot)$ corresponds to the fully connected layer that performs the gesture classification and the following softmax layer, applied to each time step $x \in \mathbf{x}$, $\text{Encoders}(\cdot)$ represents a sequence of 6 transformer encoders E , defined in the following, and PE is the Positional Encoding [17], used to encode the temporal information into the sequence. Thus, $Y(\mathbf{x})$ is a vector containing a probability distribution over n gesture classes for each time step included in \mathbf{x} .

Each transformer encoder is defined as

$$E(x) = \text{Norm}(x + \text{FC}(\text{mhAtt}(x))) \quad (4)$$

where $\text{Norm}(\cdot)$ is a normalization layer, $\text{FC}(\cdot)$ are two fully connected layers with 1024 units, followed by a ReLU activation function. The multi-head attention block mhAtt is the self-attention layer defined as

$$\text{mhAtt}(x) = (\text{Att}_1(x) \oplus \dots \oplus \text{Att}_8(x)) W^O \quad (5)$$

where

$$\text{Att}_i(x) = \text{softmax}\left(\frac{Q_i K_i}{\sqrt{d_k}}\right) V_i \quad (6)$$

Here, $Q_i = xW_i^Q$, $K_i = xW_i^K$, $V_i = xW_i^V$ are independent linear projections of x into a 64-d feature space, $d_k = 64$ is a scaling factor corresponding to the feature size of K_i , \oplus is the concatenation operator and W^O is a linear projection from and to a 512-d feature space.

4.2.3. Training

Since in the training data the large majority of frames are labelled as “no-gesture”, Group 1 uses the *Focal* loss [12] which has been shown to handle well unbalanced training datasets. In this case, they empirically verified that the recognition accuracy was higher using this loss against the standard *Categorical Cross Entropy* loss. We train the model using *Adam* [9] as optimizer with a learning rate 0.0001, weight decay 0.0001, and internal dropout set to 0.5. The features are fed to model which is trained with sliding windows of 10 time steps. We train the network on the whole training dataset for 5 epochs. The network hyper-parameters has been chosen with a k -fold cross validation (using $k = 9$).

4.2.4. Online detection

During the testing phase, the proposed system receives one frame per time. Therefore, in order to detect the beginning, the end and the class of a gesture, Group 1 implemented a *Finite State Machine* (FSM). The input is represented by a buffer, i.e. a FIFO stack, with a size of 10 frames. The FSM with 4 states, as depicted in Fig. 3 and it is running when the buffer is full of frames.

In the first state, the beginning of a gesture is detected. Each frame is classified by the proposed Transformer architecture and then, if even just a single frame is classified as gesture, the current state of the FSM is increased.

In the second state, a check on the beginning of the gesture is conducted. If not enough gestures are found in 10 consecutive windows, the FSM returns in the initial state. If at least 5 gestures are detected in 10 consecutive windows, the FSM goes in the third state. If the end of a gesture is found, the FSM passes directly on the fourth state. This last case is related to the possible presence of gestures with a very limited duration.

In the third state, the end part of the gesture is detected. The end of a gesture is an entire windows that does not contain any frame classified as gesture. In that case, the FSM passes to the fourth state.

In the fourth state, the end of the gesture is verified. Indeed, only if 25 consecutive windows do not contain any gesture, the gesture is considered completely ended. We note that this amount of windows correspond to 0.25 seconds given the acquisition frame rate of the *Leap Motion* device (25 fps). This state improves the stability of the end gesture detection. The method and the results have been computed using an Intel(R) Core(TM) i7-7700K CPU @ 4.20GHz CPU and an Nvidia GTX 1080 Ti GPU.

4.3. Group 2: Image based methods

4.3.1. Introduction

The two methods proposed by the Group 2 team are based on the transformation into images of the skeleton data captured by the Leap Motion sensor and their use for the training of a Convolutional Neural Network (CNN). The raw data points of the hand skeleton are rendered using a custom 3D visualizer: the images that constitute the dataset are captured by projecting the 3D skeleton on the xy plane. This view plane corresponds to the view from above which therefore represents the hands in a “natural” way, as a person normally sees them, and is kept constant throughout the image generation phase and at inference time. This approach has the substantial advantage of not requiring any labour-intensive feature engineering phase, so for the addition of new dynamic hand gestures to the vocabulary, the generation of the corresponding images and the retraining of the network is sufficient.

In order to provide temporal information to the network, the recent history of the gesture is represented by the fingertips traces of the hand. For the construction of the training set both the gestures of the SHREC '21 dataset and the gestures in the SFINGE3D dataset were used.

The first method used was described in detail in [1]: this method, which involves the use of a ResNet-50 for the classification of the rendered images, was retrained on the overall dataset exactly as described in the paper. In the inference phase, the same confidence thresholds used in the paper were also adopted. As shown in [1], the disadvantage of this method is that the entire temporal dimension of the gesture is condensed into the two-dimensional image in the form of fingertip traces. Although this is sufficient for the inference of some types of gestures¹, it is not for others, especially for those of typically very short duration or which develop mainly along the z axis (*grab*, *pinch*, *tap*) and therefore leave a minimal trace on the xy plane. Another disadvantage of this method is that the recognition is not continuous, but typically occurs every 300 frames (parameter to be fine-tuned depending on the dataset) to limit the number of false positives.

The second method proposed was therefore fundamentally conceived to compensate for the shortcomings of the first method: in particular, the lack of a true temporal dimension in the first method of inference led us to the training of a ResNet-3D [6], therefore able to perform 3D convolutions on volumes consisting of sequences of two-dimensional images. As for the previous

¹ The overall accuracy of the method on the SFINGE3D dataset was 0.74.

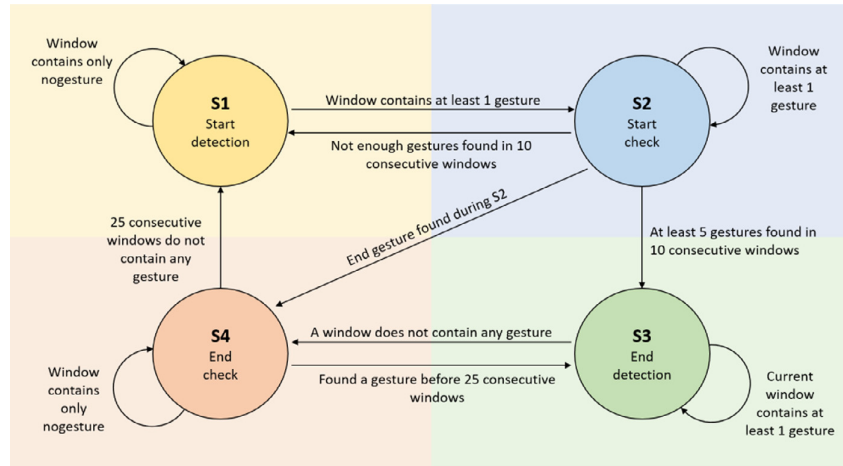


Fig. 3. Visual description of the Finite State Machine implemented for the online gesture recognition.

method, a significant offline data augmentation phase has been carried out for this method: during the generation of the training sequences, the images of the hands and fingertip traces were truncated in different moments in time to make the gestures incomplete. The continuous skeleton sequences were also sampled every 2 or 5 frames to produce the rendered images: this sampling was chosen to “compress” the longer gestures (some longer than 200 frames) into a shorter time span and as a further data augmentation. Noise was also added independently in the single rendered image, in the form of points around the skeleton of the hand to simulate residual traces of gestures.

This further data augmentation step has been added not only to help the ResNet-3D 50 converge to a robust solution² but also to try to better exploit the large model capacity of this class of networks.

This offline data augmentation phase led to the generation of 72,720 image sequences (saved in WEBP format to optimize for space) that were randomly divided with an 80%/20% split to form training and validation sets. This phase is separate from the standard online data augmentation step that occurs during network training (resize, crop, rotate, warp, brightness, contrast, saturation).

The sequence length was set to 10 frames³ and sequences shorter than *seq_len* frames were padded with black frames, while image sequences longer than *seq_len* were sampled randomly to perform further online data augmentation.

Network training was performed using the popular *Fast.ai* v2 library [7] based on Pytorch, and the progressive resizing [4] technique to optimize network convergence times. Using this technique, the training was carried out on images scaled progressively to $\frac{1}{8}, \frac{1}{6}, \frac{1}{4}, \frac{1}{3}$ of the original resolution of 1920x1080, for each training round on the network. All training rounds of the network took place in “frozen” mode, thus training only the 19-neuron (untrained) output layer of the network for 1 epoch. For the last round, the network was also trained in “unfrozen” mode, thus training all the layers of the network for a total of 7 epochs. The optimizer used was Adam and as the loss function we chose LabelSmoothingCrossEntropy. LabelSmoothingCrossEntropy is defined as:

$$L_{loss} = (1 - \epsilon)\xi(i) + \epsilon \sum \frac{\xi(j)}{N} \quad (7)$$

² This kind of network has 46.4 million parameters [10], almost double their 2D counterpart.

³ With *seq_len* = 10, the network was trained with tensors of type *torch.Size([bs, ch = 3, seq_len = 10, h, w])*

where $\xi(j)$ is cross-entropy of x and i is the correct class. Through LabelSmoothingCrossEntropy we try to compensate for noisy labels in the training set: instead of wanting the model to predict 1 for the correct class and 0 for all the others, we teach it to predict $1 - \epsilon$ for the correct class and ϵ for all the others, being ϵ a small positive constant and N the number of classes of the problem.

The training took place on a GPU node of the new high-performance EOS cluster located within the University of Pavia. This node has a dual Intel Xeon Gold 6130 processor (16 cores, 32 threads each) with 128 GB RAM and 2 Nvidia V100 GPUs with 32 GB VRAM each. The training took place using PyTorch’s Data-Parallel mode, so the batch size for the different training rounds was set to 32, 16, 8, 4 respectively to best occupy all the 64 GB of VRAM available on the two GPUs. The learning rate was set to $1e-2, 1e-2, 1e-2, 1e-4$ respectively for the “frozen” training rounds and $4e-4$ for the 7 “unfrozen” epochs in the last round of training. The final accuracy of the model used for the submission was 0.963 against the validation set.

Unfortunately, although these methods are promising in terms of approach to the problem, the short duration of the contest did not allow the team to optimize the results. As for the first method, it would probably have been sufficient to slightly lower the confidence threshold to consider a gesture as recognized to raise the positive detection score (possibly to the detriment of the false positive score which is still sufficiently low). The main problem of the second method, on the other hand, is certainly the too short sequence length: during the training the spatial resolution was privileged to the detriment of the temporal one and this did not allow the network to have sufficient context to learn. It is necessary to explore the trade-offs between spatial and temporal resolution to verify under which conditions the network learns best.

4.4. Group 3: uDeepGRU and TSGR

4.4.1. Introduction

For this track, Group 3 submitted results obtained from two different methods: the improved uDeepGRU model [2,13] as well as a novel method dubbed *Temporal Shift Gesture Recognizer* (TSGR). Both of their methods are based on deep neural networks. In the following they provide an overview of each one.

4.4.2. uDeepGRU

Group 3 introduced the improved version of uDeepGRU [13]. Fig. 4 depicts the network architecture of the improved model.

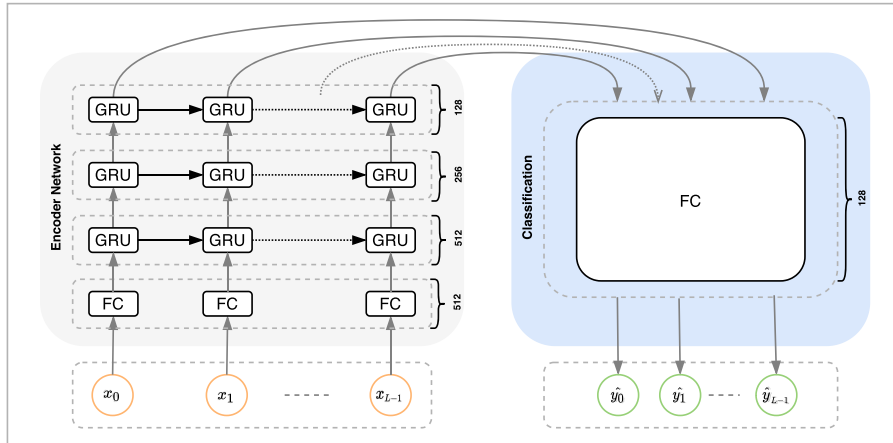


Fig. 4. The improved uDeepGRU architecture which consists of an encoder network and a classification subnetwork.

The uDeepGRU model is based on recurrent neural networks (RNN) and uses unidirectional gated recurrent units (GRU) as its main building block. Frames of a gesture sequence are sequentially fed to the network, and the network outputs the predicted label for every frame. Concretely, the network takes as input the feature vector x_t at time step t and produces the output label $\hat{y}_t \in \{\text{None}\} \cup C$, where C is the set of all possible gestures in the dictionary and *None* indicates a no gesture. The the transition equation for each GRU cell in the uDeepGRU model is defined as:

$$\begin{aligned}
 r_t &= \sigma \left((W_x^r x_t + b_x^r) + (W_h^r h_{(t-1)} + b_h^r) \right) \\
 u_t &= \sigma \left((W_x^u x_t + b_x^u) + (W_h^u h_{(t-1)} + b_h^u) \right) \\
 c_t &= \tanh \left((W_x^c x_t + b_x^c) + r_t (W_h^c h_{(t-1)} + b_h^c) \right) \\
 h_t &= u_t \circ h_{(t-1)} + (1 - u_t) \circ c_t
 \end{aligned} \tag{8}$$

where σ is the sigmoid function, \circ denotes the Hadamard product, r_t , u_t and c_t are reset, update and candidate gates respectively and W_p^q and b_p^q are the trainable weights and biases. The initial hidden state tensor h_0 of all the GRUs in our model is initialized to zero.

The improved uDeepGRU model introduces two main changes compared to the original implementation in [2]. First, there is an extra feature extraction layer at the beginning of the model. This layer consists of a fully-connected (FC) layer with $\tanh()$ activation. Given x_t , the feature vector of an input frame at time step t , this layer computes $f_t = \tanh(x_t W^T + b)$, where W and b are trainable weights and the bias term respectively. The goal of this feature extraction layer is to increase the feature extraction capacity of the model. Second, the training objective function for this model is the focal loss function [12]. This loss function was originally proposed for problems with unbalanced labeled data. The online gesture recognition problem is an example of such unbalanced data, as many frames obtained from an input device are typically non-gestural interactions. Focal loss attempts to dynamically weight the standard cross-entropy loss to dampen the effect of well-classified examples on the final loss value. For a model which outputs p as the probability of a given class label, the focal loss is defined as:

$$L_{\text{focal}} = -(1 - p)^\gamma \log(p) \tag{9}$$

where $\gamma \geq 0$ determines how much emphasis is put on misclassified examples. We use $\gamma = 1$ in our implementation and train uDeepGRU end-to-end.

4.4.3. TSGR

Although RNN-based models have demonstrated great capabilities in sequence modeling and prediction tasks, they suffer from a few critical drawbacks. Namely, RNNs rely on their hidden state tensors for modeling the relationships across different time steps of their inputs. Also, most RNN-based models suffer from problems such as vanishing or exploding gradients which make them difficult to train.

To address these issues, Lin et al. [11] recently introduced the *Temporal Shift Module* (TSM) an alternative sequence modeling paradigm using the temporal shifting of input features. The basic idea behind such methods is simple, yet powerful: at each time step during the processing of a temporal sequence, replace a portion of the features of the current time step with those of other time steps via shifting those features across the temporal dimension. This allows the network to perform temporal modeling across extracted features. Depending on the task, the shifting operation can either be bidirectional or unidirectional. Unidirectional shifting is suitable for online recognition tasks as during recognition, only the features of the prior time steps are available.

Based on this idea Group 3 devised TSGR, their second online gesture recognizer. The TSGR model is depicted in Fig. 5. Similar to uDeepGRU, our TSGR model takes the input features for every time step and produces the output class prediction. This model, which is conceptually simpler than uDeepGRU, consists of only FC layers with TSM layers in between. Each TSM layer replaces half of the features of the current frame with those of the past 5th frame. This implies that the amount of feature shifting in our model is five frames. Although the original TSM model [11] shifts the features by one, Group 3 experimentally found five shifts to work better for this track's data.

We refer to the combination of FC and TSM layers as *shift nodes* (SN) henceforth. Each SN consists of a TSM layer along with two FC layers, namely FC_{Shift} and FC_{Residual} with $ReLU()$ activations⁴. The dimensionality of both FC layers is the same, however FC_{Residual} layers do not include a bias term. The sequence of operations inside each node is as follows. At each time step t , the node takes the feature vector f_t , performs the shifting operation by replacing half of the values in f_t by those computed for f_{t-5} . The node then passes these features through FC_{Shift} and saves the intermediate results f_t^{Shift} . The node also computes and stores the intermediate result $f_t^{\text{Residual}} = FC_{\text{Residual}}(f_t)$. The final output of the node is computed as $f_t^{\text{SN}} = f_t^{\text{Shift}} + f_t^{\text{Residual}}$. Additionally batch normal-

⁴ Except for the very first layer, where $\tanh()$ is used

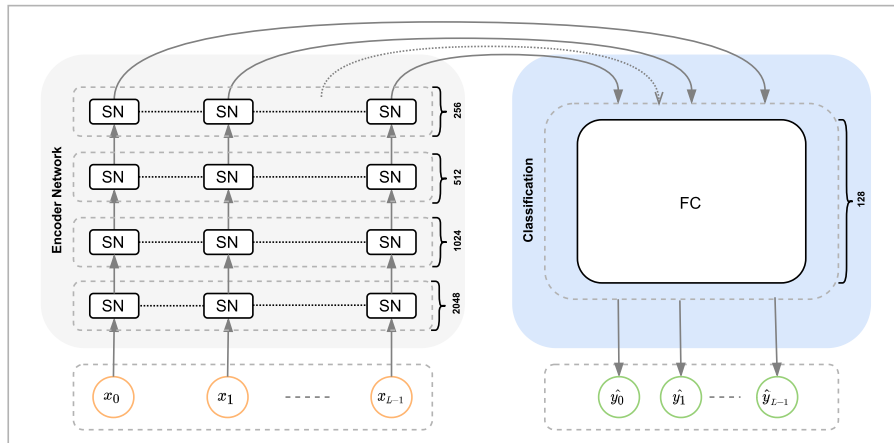


Fig. 5. The TSGR architecture which consists of an encoder network and a classification subnetwork. Each shift node (SN) has access only to the features of the previous time steps in that same layer.

ization [8] is applied before the final result is passed to the next layer. Similar to uDeepGRU, we optimize the focal loss function to train TSGR end-to-end.

4.4.4. Implementation and training

Each frame of the data is treated as one 60-dimensional vector x_t obtained by concatenating the 3D position of all joints. Every feature vector was z-score normalized using the mean and standard deviation of all feature vectors in the training set.

Group 3 implemented their models in PyTorch. Both models are trained end-to-end on the training set, with six random sequences withheld for validation. Training was done using the Adam optimizer [9] with a learning rate of 0.0002 and a mini-batch size of 10. The maximum length of a training sample was fixed to 256 (longer samples were split to chunks of at most 256 frames). Counter-intuitively, they found any kind of data augmentation on the training set to be harmful. After training concluded, we saved the model that produced the best F_1 score on the validation set.

4.4.5. Results

At test time, Group 3 runs each test sample through the network and obtain per-frame class labels \hat{y}_t . They do not perform any post-processing on the output results and the hardware used for the compute the results is a AMD Ryzen 3900x processor and an NVIDIA Titan RTX GPU with 24 GB VRAM.

They obtained three sets of results. The first set consists of the results of an ensemble of 10 uDeepGRU models, each trained on a different portion of the training set. The second set is the classification labels obtained from a single TSGR trained model, and the last set contains the results of an ensemble of 15 TSGR models, each trained on a different portion of the training set.

4.5. Group 4: Spatial-Temporal Graph Convolutional Networks

4.5.1. Data preprocessing

Based on the task description, the dataset is mixed with random hand movements between real gestures. Group 4 added those noisy segments labeled as non-gestures into the training procedure to improve robustness of the model. Next, they extract both gestures and non-gestures data with a length of 200 frames for each segment. Due to the limitation on training data, a stratified 5-fold based on class-distribution is applied to avoid under and over fitting.

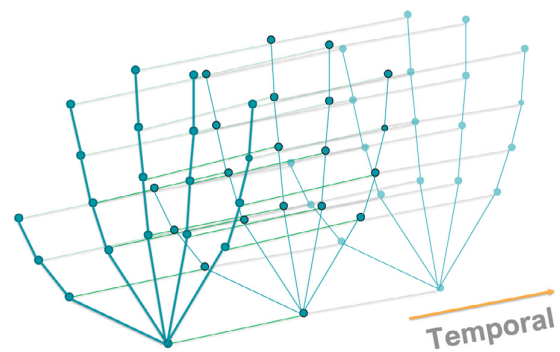


Fig. 6. Spatial-temporal graph neural networks for hand gesture recognition.

4.5.2. Classification model: Spatial-Temporal Graph Convolutional Networks

Spatial-Temporal Graph Convolutional Networks (ST-GCN) [20] is an extended version of graph neural networks to a spatial-temporal graph model. The graph can learn patterns embedded in the spatial configuration by exploring locality of graph convolution as well as temporal dynamics. As proposed by Yan et al., Group 4 constructs a sequence of skeleton graphs, each node represents a joint of the hand. Moreover, there are spatial edges for building up the connectivity of joints according to natural structure of human hands and temporal edges connecting the same joints across continuous frames of actions, described as Fig. 6. Also, each node has its own features composed of 3D coordinates and quaternions. In the classification module using ST-GCN, features of a joint at frame t is a vector with length 7 consist of $[x, y, z, a, b, c, d]$ where (x, y, z) expresses 3D coordinates and a quaternion q is described by $q = a + bi + cj + dk$.

In the detection module, we only use (x, y, z) to indicate potential candidates.

4.5.3. Energy-based detection module

With respect to detection and localization module, Group 4 uses a non-deep learning approach by leveraging the shift of every joint of human hand over sequences of consecutive frames. They define an energy-based function to calculate the amount of energy

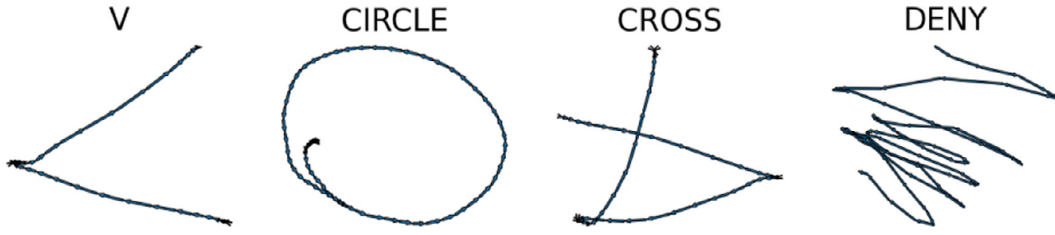


Fig. 7. 2D Trajectories of hand gestures.

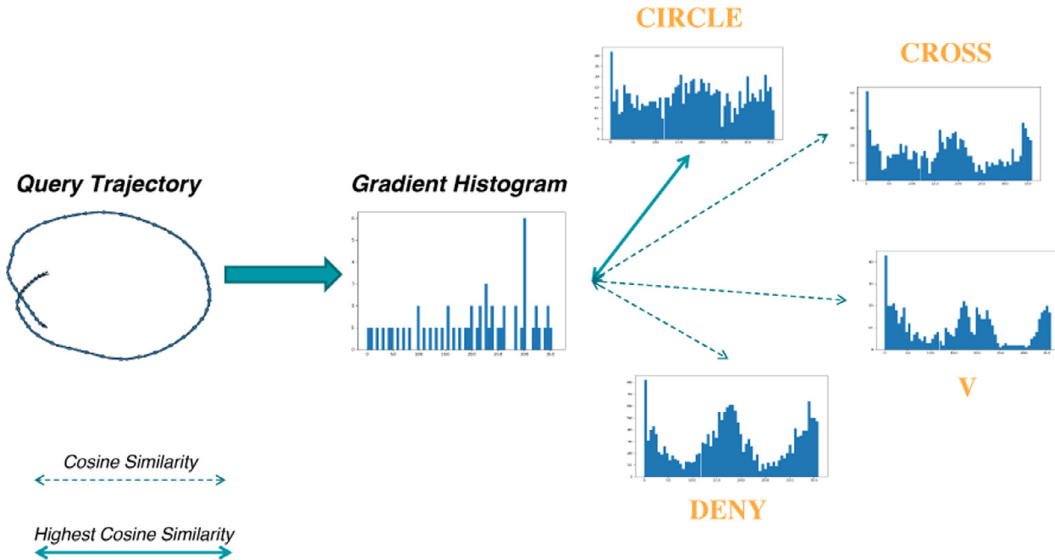


Fig. 8. Cosine similarity between gradient histograms.

accumulated in a window of length L inside the gesture stream:

$$E(w) = \sum_{n=1}^N \sum_{t=1}^L \sqrt{\left(\frac{w_{n,t}^x}{w_{n,t-1}^x} - 1\right)^2 + \left(\frac{w_{n,t}^y}{w_{n,t-1}^y} - 1\right)^2 + \left(\frac{w_{n,t}^z}{w_{n,t-1}^z} - 1\right)^2} \quad (10)$$

Index t ranges from 1 to L . n is the index of the hand joint ranging from 1 to N . So, $w_{n,t}^x$ represents the x-coordinate of n^{th} joint at frame t .

Using Eq. 10, Group 4 adopted a sliding window approach, estimating the value of $E(w_i)$ on multiple windows w_i of length L starting at different locations, with consecutive windows w_i, w_{i+1} separated by a fixed stride step.

Authors then take as candidate gestures all the windows corresponding to local maxima of the energy, e.g. those satisfying the following conditions:

$$\begin{aligned} \delta E(w_i) &< \epsilon, \\ \delta E(w_{i-1}) &> 0, \\ \delta E(w_{i+1}) &< 0 \end{aligned} \quad (11)$$

$\delta E(w_i)$ is the time derivative of energy of segment w_i , approximated by:

$$\delta E(w_i) = \frac{E(w_{i+1}) - E(w_{i-1})}{(i+1) - (i-1)} \quad (12)$$

After detecting possible segments, Group 4 feeds those candidates into the ST-GCN model. Those segments that are predicted to belong to a non-gesture class with confidence score $>$ threshold α or be a member of gesture classes with confidence score $<$

threshold β are filtered. Threshold α and β are chosen based on the validation data.

4.5.4. Trajectory-based fine-tuning with PCA and gradient histogram

Group 4 noticed that some gestures can be easily classified by leveraging their trajectory. Among the 20 joints of a human hand in the dataset, they chose IndexEnd of an index finger to consider action's trajectory for gestures: CIRCLE, V, CROSS, DENY. The principal component analysis is applied to reduce 3D coordinate to 2D system, see Fig. 7.

Next, those sequences of 2D coordinates are utilized to find gradient vectors with O_x and O_y axis and their angles. A histogram with N bins ranging from $-\pi/2$ to $\pi/2$ is computed on the frequency of calculated angles [15]. This 1D orientation histogram is viewed as a feature vector of each gesture. With regard to the training set, every class of gestures comprises some segment members that belong to this class. Thus, each class of gesture is represented by the mean histogram feature vector of its members. In the inference phase, they extract the gradient histogram feature vector of every candidate segment and then compare it with each class's representation vector using Cosine Similarity, described as Fig. 8. Finally, the highest similarity score is ensemble with confidence score from the ST-GCN model and the predicted label is returned.

4.5.5. Experiments

Group 4 did experiments on single ST-GCN models, k-fold models, and ensemble models as well. Finally, they chose 3 configurations of them that correspond to 3 RUNS.

- RUN1: Single ST-GCN model.

Table 1

Average scores obtained by all the proposed techniques on the test sequences and corresponding execution times for the classification of all the test data and for a single prediction at a given sequence point (online detection). Bold fonts indicate best results.

Method	Det. Rate	FP Rate	Jac. Ind.	Tot.Time(s)	Class.Time(s)
Baseline	0.3993	0.7639	0.2566	1161.0	0.08
Group 1 - Run 1	0.7014	0.3576	0.5501	435.5	1.36
Group 1 - Run 2	0.6875	0.5521	0.4737	435.5	1.36
Group 1 - Run 3	0.7292	0.2569	0.6029	435.5	1.36
Group 2 - Run 1	0.4861	0.9271	0.2772	48781.0	0.41
Group 2 - Run 2	0.4931	0.1667	0.4458	4897.2	0.81
Group 3 - Run 1	0.6042	0.3021	0.4987	66.7	0.6×10^{-4}
Group 3 - Run 2	0.7569	0.3403	0.6194	0.4	0.3×10^{-5}
Group 3 - Run 3	0.7431	0.2708	0.6238	5.1	0.4×10^{-4}
Group 4 - Run 1	0.8403	0.0903	0.7925	94.6	0.6×10^{-2}
Group 4 - Run 2	0.8750	0.0556	0.8353	281.4	0.03
Group 4 - Run 3	0.8993	0.066	0.8526	289.2	0.16

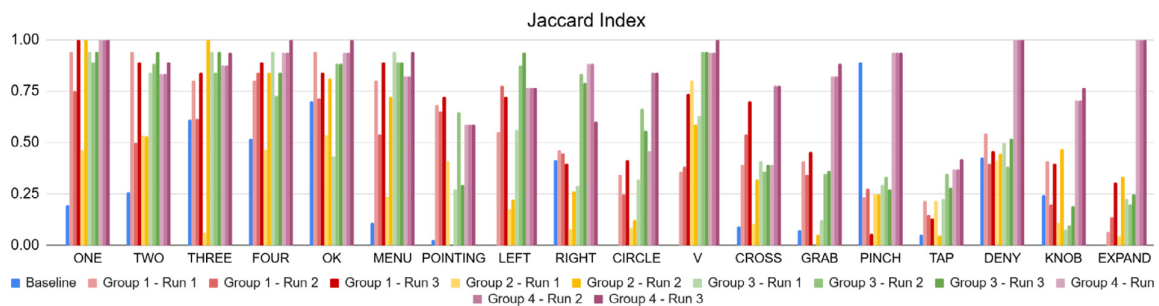


Fig. 9. Jaccard index per class on all the test sequences.

(CPU Intel core i5-8300H 2.3GHz with Turbo Boost up to 4.0GHz)

- RUN2: ST-GCN model with stratified 5-fold. (GPU Tesla P100 16GB)
- RUN3: 5-fold ST-GCN models ensemble with gradient histogram. (CPU Intel Core i5-8300H 2.3GHz with Turbo Boost up to 4.0GHz)

5. Evaluation results

A summary of the results for each group, averaged over all the gestures, is presented in Table 1. The table also shows the results of the execution time of the methods by reporting the Total Time and the Classification Time. The first is a measure of the time each method takes to compute results for the entire test set (i.e. all the sequences), the second, is a measure of the average time needed, for each method, to perform a single gesture classification. Times show that all the methods are suitable for real-time applications, even if the performances are not directly comparable as the software has been executed on different architectures.

The methods based on ST-GCN provide clearly the best performance, and it is interesting also to note that the combination of the network based method with simple heuristics is able to improve the scores. These methods also provide a low number of false positives, that are instead not negligible in the other methods.

To better understand the outcomes of the different techniques it is useful to analyze the scores related to the single gesture classes.

The bar charts in Figs. 9, 10 and 11 show the per-class scores of all the methods.

As expected, all the methods present good performances for the static gestures (i.e. from ONE to POINTING). The average JI for these gestures is 0.73. However, some methods result in a non negligible number of false positives that could make difficult to use the tech-

niques in a practical scenario. The average false detection rate on static gestures is, in fact, 24%, meaning a false positive detected every 4 gesture recognized, that is quite high for practical purposes.

The POINTING gesture is the most challenging static one, and it is reasonable, being the one with more relevant variations in the execution.

Dynamic gestures are much harder to recognize (average JI 0.46, average FD 0.41), with the notable exception of the ST-GCN based methods. As shown in Fig. 12, considering the best run (i.e. highest Jaccard Index) for each group, the results show that among the dynamic gestures, those considered fine and therefore characterized also by the single fingers' trajectories (GRAB, PINCH, TAP, DENY, KNOB, EXPAND), present further issues for most of the methods. ST-GCN works well on most of the gesture class with just a few exceptions (TAP, POINTING, KNOB).

Some surprising facts appear from the analysis of single classes: for example the LEFT gesture presents a very low number of false positives despite being a dynamic one, while the RIGHT gesture is found in a lot of false detections.

The PINCH gesture is surprisingly detected easily and with few false detections by the baseline method, and it is hardly detected by most of the network based techniques.

6. Discussion

The evaluation outcomes provide useful insights for the design of online gesture recognizers usable "in the wild". The techniques tested provide promising scores given a limited number of annotated sequences for training. Given the short amount of time available for the contest this is a good result.

A nice aspect of the submissions received is that the proposed methods are quite different from each other and are exemplars of the principal network-based approaches proposed in the literature for these tasks, namely Recurrent Networks, Graph Networks, Transformer Networks, and Convolutional Neural Networks. The fact that Spatiotemporal Graph Convolutional networks provide the



Fig. 10. Detection rate per class on all the test sequences.

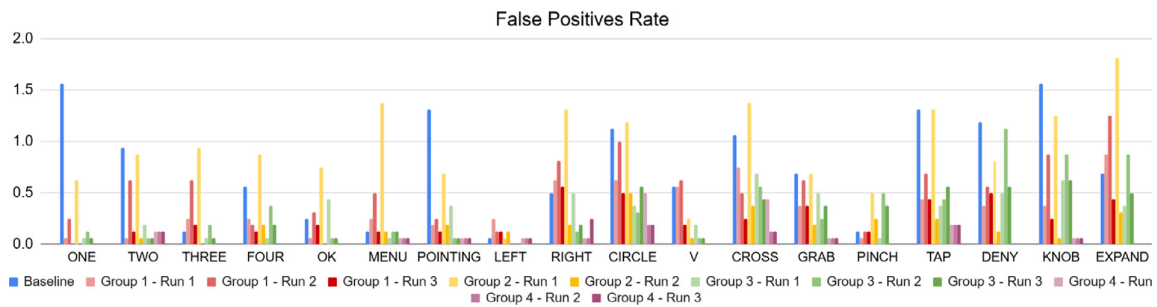


Fig. 11. False Positives rate per class on all the test sequences.

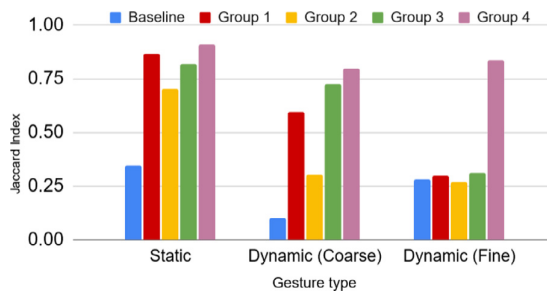


Fig. 12. Jaccard Index by gesture type (best run for each group).

best results is consistent with the literature on action recognition from body skeletons, where the best scores on the related benchmarks have been obtained with similar approaches. However, on hand gesture recognition, good results on old benchmarks have been obtained with modified versions of LSTM including neighbor information [14] or using 1D CNN [21]. These methods have not been proposed in this contest, and we plan to test them on our data as future work.

It must be noted, then, that the performances of the different techniques depend on hyperparameters tuning and training data augmentation, and, given the short amount of time available for the contest, the results could be improved and the ranking changed.

In any case, it is worth noting that for all the proposed methods but the ST-GCN-based dynamic gestures are not well handled and false positives are a relevant issue.

With all the methods, selected gestures were hard to be detected (e.g. TAP, KNOB, POINTING). A possible future research direction is therefore to investigate the reasons for these problems, which may rely on similarities between segments of different gestures or the variability in the execution. The goal could be to create optimal dictionaries for gestural interfaces avoiding the inclusion of "problematic" classes.

The improvements of the ST-GCN results obtained with the addition of simple handcrafted similarity evaluations and the fact that simple handcrafted features work well on specific gestures (PINCH) show that it is hard to have a generic method well-suited for the recognition of all the gesture types. A viable option to address this issue could be to define multiple recognizers for specific gestures.

Problems with dynamic gestures could also derive from the limited numbers of training sequences and subjects performing the gestures and from the fact that the users performing the gestures in the test set were not involved in the recording of the training set. The availability of larger and more varied training sets could be exploited to increase the detection performances. However, for the practical use of recognizers in interface design, the availability of gesture recognizers that can be trained with few examples would be particularly useful.

Another aspect that should be investigated is related to the computational load required for the online classification. While all the networks can be used for online recognition on a high-end PC, the possibility to have them running on HoloLens or Oculus Quest needs to be checked.

We plan to update the dataset by adding new data. Furthermore, while the current gestures were recorded within a continuous gesticulation, but separated by non-gesture actions, we want to record sequences with series of adjacent gestures and design a new task involving the detection of series of atomic gestures.

7. Conclusions

The development of effective and flexible gesture recognizers able to detect and correctly classify hand gestures of different kinds is fundamental not only to enable advanced user interfaces for Virtual and Mixed Reality applications, but also, for example, to enable the realization of touchless interfaces like public kiosks, that are expected to replace touch-based ones after the emergence of the pandemic issues, being a more hygienic and safer solution. It is, therefore, important to support the research on this kind of

tool, developing benchmarks able to test the algorithms on realistic user scenarios. The SHREC 2021: Track on Skeleton-based Hand Gesture Recognition in the Wild tries to do this. We believe that the dataset created and the methods proposed by the participants will be a useful asset for the researchers working on this topic.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRedit authorship contribution statement

Ariel Caputo: Methodology, Software, Validation, Formal analysis, Writing – original draft, Writing – review & editing. **Andrea Giachetti:** Methodology, Software, Validation, Formal analysis, Writing – original draft, Writing – review & editing. **Simone Soso:** Methodology, Software. **Deborah Pintani:** Methodology, Software. **Andrea D'Eusanio:** Methodology, Software, Writing – original draft. **Stefano Pini:** Methodology, Software, Writing – original draft. **Guido Borghi:** Methodology, Software, Writing – original draft. **Alessandro Simoni:** Methodology, Software, Writing – original draft. **Roberto Vezzani:** Methodology, Software, Writing – original draft. **Rita Cucchiara:** Methodology, Software, Writing – original draft. **Andrea Ranieri:** Methodology, Software, Writing – original draft. **Franca Giannini:** Methodology, Software, Writing – original draft. **Katia Lupinetti:** Methodology, Software, Writing – original draft. **Marina Monti:** Methodology, Software, Writing – original draft. **Mehran Maghoumi:** Methodology, Software, Writing – original draft. **Joseph J. LaViola Jr:** Methodology, Software, Writing – original draft. **Minh-Quan Le:** Methodology, Software, Writing – original draft. **Hai-Dang Nguyen:** Methodology, Software, Writing – original draft. **Minh-Triet Tran:** Methodology, Software, Writing – original draft.

References

- [1] Caputo A, Giachetti A, Giannini F, Lupinetti K, Monti M, Pegoraro M, Ranieri A. Sfinge 3d: a novel benchmark for online detection and recognition of heterogeneous hand gestures from 3d fingers' trajectories. *Computers & Graphics* 2020;91:232–42. doi:10.1016/j.cag.2020.07.014. <https://www.sciencedirect.com/science/article/pii/S0097849320301163>
- [2] Caputo FM, Burato S, Pavan G, Voillemin T, Wannous H, Vandeborre JP, Maghoumi M, Taranta II EM, Razmjoo A, LaViola Jr JJ, Manganaro F, Pini S, Borghi G, Vezzani R, Cucchiara R, Nguyen H, Tran MT, Giachetti A. Online Gesture Recognition. In: *Eurographics Workshop on 3D Object Retrieval*; 2019.
- [3] De Smedt Q, Wannous H, Vandeborre J-P, Guerry J, Le Saux B, Filliat D. Shrec'17 track: 3d hand gesture recognition using a depth and skeletal dataset. In: *3DOR-10th Eurographics Workshop on 3D Object Retrieval*; 2017. p. 1–6.
- [4] Fang L, Monroe F, Novak SW, Kirk L, Schiavon CR, Yu SB, Zhang T, Wu M, Kastner K, Kubota Y, Zhang Z, Pekurnaz G, Mendenhall J, Harris K, Howard J, Manor U. Deep learning-based point-scanning super-resolution imaging. *bioRxiv* 2019. doi:10.1101/740548.
- [5] Han S, Liu B, Cabezas R, Twigg CD, Zhang P, Petkau J, Yu T-H, Tai C-J, Akbay M, Wang Z, et al. Megatrack: monochrome egocentric articulated hand-tracking for virtual reality. *ACM Transactions on Graphics (TOG)* 2020;39(4): 87–1
- [6] Hara K, Kataoka H, Satoh Y. Can spatiotemporal 3d CNNs retrace the history of 2d CNNs and imagenet?, *arXiv:1711.09577*.
- [7] Howard J, Gugger S. Fastai: a layered API for deep learning. *Information* 2020;11(2):108. doi:10.3390/info11020108.
- [8] Ioffe S, Szegedy C. Batch normalization: Accelerating deep network training by reducing internal covariate shift; 2015. p. 448–56. URL <http://jmlr.org/proceedings/papers/v37/loff15.pdf>
- [9] Kingma DP, Ba J. Adam: a method for stochastic optimization, *arXiv preprint arXiv:1412.6980*.
- [10] Leong MC, Prasad DK, Lee YT, Lin F. Semi-CNN architecture for effective spatio-temporal learning in action recognition. *Applied Sciences* 2020;10(2):557.
- [11] Lin J, Gan C, Han S. Tsm: Temporal shift module for efficient video understanding. In: *Proceedings of the IEEE International Conference on Computer Vision*; 2019.
- [12] Lin T-Y, Goyal P, Girshick R, He K, Dollár P. Focal loss for dense object detection. In: *Proceedings of the IEEE international conference on computer vision*; 2017. p. 2980–8.
- [13] Maghoumi M. Deep Recurrent Networks for Gesture Recognition and Synthesis. University of Central Florida Orlando, Florida; 2020. Ph.D. thesis.
- [14] Min Y, Zhang Y, Chai X, Chen X. An efficient pointlstm for point clouds based gesture recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; 2020. p. 5761–70.
- [15] Thompson EM, Biasotti S, Giachetti A, Tortorici C, Werghi N, Obeid AS, Berretti S, Nguyen-Dinh H-P, Le M-Q, Nguyen H-D, Tran M-T, Gigli L, Velasco-Forero S, Marcotegui B, Sipiran I, Bustos B, Romanelis I, Fotis V, Arvanitis G, Moustakas K, Otu E, Zwigelaar R, Hunter D, Liu Y, Arteaga Y, Luxman R. SHREC'20 Track: retrieval of digital surfaces with similar geometric reliefs. *Computers and Graphics* 2020;91:199–218. doi:10.1016/j.cag.2020.07.011. URL <https://hal.archives-ouvertes.fr/hal-02916788>
- [16] Ungureanu D, Bogo F, Galliani S, Sama P, Duan X, Meekhof C, Stühmer J, Cashman TJ, Tekin B, Schönberger JL, et al. Hololens 2 research mode as a tool for computer vision research, *arXiv preprint arXiv:2008.11239*.
- [17] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L, Polosukhin I. Attention is all you need. In: *Advances in neural information processing systems (NIPS)*; 2017. p. 5998–6008.
- [18] Wan J, Zhao Y, Zhou S, Guyon I, Escalera S, Li SZ. Chalearn looking at people rgb-d isolated and continuous datasets for gesture recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*; 2016. p. 56–64.
- [19] Weichert F, Bachmann D, Rudak B, Fisseler D. Analysis of the accuracy and robustness of the leap motion controller. *Sensors* 2013;13(5):6380–93.
- [20] Yan S, Xiong Y, Lin D. Spatial temporal graph convolutional networks for skeleton-based action recognition. *CoRR* 2018. arXiv:1801.07455.
- [21] Yang F, Wu Y, Sakti S, Nakamura S. Make skeleton-based action recognition model smaller, faster and better. In: *Proceedings of the ACM Multimedia Asia*; 2019. p. 1–6.
- [22] Zhang F, Bazarevsky V, Vakunov A, Tkachenka A, Sung G, Chang C-L, Grundmann M. Mediapipe hands: on-device real-time hand tracking, *arXiv preprint arXiv:2006.10214*.
- [23] Zhang Y, Cao C, Cheng J, Lu H. Egogesture: a new dataset and benchmark for egocentric hand gesture recognition. *IEEE Trans Multimedia* 2018;20(5):1038–50.