

2D Pose Estimation based Child Action Recognition

Sanka Mohottala, Sandun Abeygunawardana, Pradeepa Samarasinghe, Dharshana Kasthurirathna
Faculty of Computing, Sri Lanka Institute of Information Technology, Sri Lanka
sanka.m@sliit.lk, it19362854@my.sliit.lk, pradeepa.s@sliit.lk, dharshana.k@sliit.lk

Charith Abhayaratne

Department of Electronic and Electrical Engineering, University of Sheffield, United Kingdom
c.abhayaratne@sheffield.ac.uk

Abstract—We present a graph convolutional network with 2D pose estimation for the first time on child action recognition task achieving on par results with LRCN on a benchmark dataset containing unconstrained environment based videos.

Index Terms—child action recognition, graph convolutional networks, long-term recurrent convolutional network, transfer learning

I. INTRODUCTION

Human activity recognition (HAR) has been a focused research area due to its diverse applications in human-computer interaction [1], surveillance [2] and health care [3]. The main goal of HAR is to identify actions performed by one or more humans in a temporal sequence of observations. The increased computational power together with the availability of public datasets have enabled the training of large networks, such as multi-stream 3D Convolutional Neural Network (CNN) architectures [4], significantly boosting action recognition performance.

Child action recognition (CAR) has important applications in safety monitoring [5], development assessment [6], and others. However, children are largely underrepresented in both pose estimation datasets [7] and HAR datasets. This can be explained by the fact that data from adult activities can be used in many more applications than data from children activities. Moreover, some of the largest annotated human activity datasets rely on user uploaded videos to public platforms [8], [9], in which the presence of children is limited due to privacy concerns.

RGB modality being one of the oldest and most used method of acquiring motion information, many deep learning techniques on HAR have achieved best accuracy on RGB benchmark datasets like UCF101 and Kinetics600 [8], [10]. Among those state-of-the-art methods, Long-Term Recurrent Convolutional Network (LRCN) architecture was selected as the best approach in this study for the HAR [11] due to its capability in handling variable length input and learning complex video sequences.

While RGB modality has been the prominent approach for HAR, different non-RGB approaches have also come into play in the past decade with the revival of deep learning. Veeriah et al. [12] introduced a modified LSTM approach where HOG3D was used to extract features from the image and they were used as the input to the model. Fernando et al. [13] introduced novel

feature extraction methods and a ranking machine approach to learn a hierarchical representation of actions from video frames. These approaches have been used with HAR datasets such as Kinetics-400 [14] and have achieved comparable results [15], [16].

Human pose estimation through 2D keypoint estimation is another feature extraction approach that has been popularised recently due to emergence of libraries such as OpenPose [17], and BlazePose [18] where the resultant features resemble the human skeleton, often referred as the skeleton modality. While skeleton modality approaches with LSTM [19] and temporal convolution models (TCN) [20] have outperformed other non-RGB approaches, graph convolution (GCN) approaches such as ST-GCN [16], 2s-AGCN [15] have surpassed both these approaches. ST-GCN being the first GCN implementation for HAR, in this research, we use it for skeleton modality implementations. OpenPose was used for pose estimation since [7] shows that in child pose estimation, OpenPose performs better than other methods even when truncations and occlusions are present.

Most of the CAR research done with skeleton modality are based on RGB+D datasets but there are few using RGB datasets and pose estimation techniques. An LSTM based model for stereotypical action recognition of children for ASD detection is introduced in [21], and [22], [23] recognized gross motor actions of 4-5 years old children with OpenPose as the pose estimator. Authors in [21] have introduced methods to overcome truncation and occlusion issues resulting in an overall 90% accuracy. For tracking, they have utilized the distance of corresponding skeleton joints in frame sequence. A CNN approach has been used in [22] achieving 82% accuracy, where tracking was done using a particle filter algorithm and skeleton standardization was applied such that the output looks as if the camera angle and skeleton size remain the same. Improving this further the authors have introduced a model similar to TCN in [23], where an improved standardization and a novel data augmentation were utilized to achieve $\approx 99\%$ accuracy. All these CAR implementations in skeleton modality have used data captured in constrained environments, hence they work with stabilized skeleton sequences. Since our interest is in unconstrained environments, we are using annotated subsets from Kinetics-400 and Kinetics-600 datasets in this research.

Most of the datasets used in CAR research are either private

datasets as in [21]–[23] and the openly available ones such as [24], [25] contain only the skeleton data. While there are handful of datasets available such as [26], [27] they contain only a small number of videos. With the release of our annotated dataset along with the benchmark results detailed in the results section, we provide opportunities for further research on annotated public datasets.

As evidenced through the inferior performance of C3D and Resnet50+LSTM models in [28], the skeleton modality has outperformed RGB modality on datasets created in constrained environments [19]. But its performance on unconstrained environment based datasets [14] is inferior to RGB modality as shown by ST-GCN and PoseC3D models in [16], [29]. While these datasets are mostly adult based datasets, [16] hypothesize that this performance degradation is due to information loss of object/scene interaction and shows through a Kinetics-400 subset called Kinetics-motion that when the activities are highly motion oriented, skeleton modality achieves performance similar to RGB modality. Motivated by this insight, in this research we attempt to achieve comparable results on unconstrained skeleton modality child activity datasets.

With the changes done to the model architecture, OpenPose V2 (2019) has improved accuracy by 7% compared to OpenPose V1 (2016) [17]. OpenPose outputs the locations of skeleton joint along with a confidence value that give a measure about the reliability of the inferred joint. We attempt to see if there is any correlation between accuracy of models and the average confidence values.

Based on the above literature study and the gaps identified, in this paper we provide the following main contributions.

- To the best of our knowledge this is the first GCN based CAR implementation done on extracted 2D skeleton sequences from unconstrained environment videos
- We show that when pre-trained with large datasets, ST-GCN on skeleton modality achieves on par accuracy with LRCN on RGB modality.
- We show there is no strong correlation between confidence value of skeleton sequences and the class-wise accuracy of the model.
- We provide the annotated dataset along with detailed guideline to continue further research¹.

The rest of the paper is organized as follows. The implemented model architectures and datasets are discussed in Section II. Experimental setups used in each case and the evaluation methods are discussed in Section III. The performance of different approaches and comparison of models are discussed in Section IV while Section V concludes with future research directions.

II. METHODOLOGY

A. Datasets

Due to the scarcity of public child datasets, an annotated child dataset was created using an eight class subset of

Kinetics-600 dataset. Based on the appearance of the main performer of the action, each video was labeled as child or adult. Kinetics-600 dataset was chosen since the activities are already classified and time stamped and since the videos are taken from YouTube, they belong to real-world scenarios. Eight classes that were selected are detailed in Table I. They were selected based on the similarity to Child-Whole Body Gesture (CWBG) dataset [24]. Majority of selected classes are motion oriented actions since we are interested in such actions as detailed in Section I.

1) *Kinetics-600 Subset (KS)*: Kinetics-600 contains $\approx 480,000$ (mostly-adult) activity videos with an average duration of 10 seconds taken from YouTube. While the dataset contains 600 different activities with 600 videos per class on average, they vary from atomic actions like "Squat" to hierarchical activities like "Playing poker". Composition of the annotated kinetics-600-Subset is given in Table I. Four protocols were introduced for the model development and random 75% train set, 25% test set splitting was done.

- KS-Full: Child data of all eight classes included.
- KS-Large: Only the five classes with highest child data percentage are included (i.e. first five rows of Table I).
- KS-Balanced: Child data of same five classes as above but random sampling of 250 videos per class was done to create a balanced dataset.
- KS-Small-C: Only child data of three classes with lowest child data percentage are included (i.e. last three rows of Table I).
- KS-Small-A: Same three classes as in KS-Small-C, but contains only adult data.

TABLE I: Child-Adult video distribution of KS

Class Name	Child Data	Adult Data	Child Percentage
Hopscotch	643	135	83%
Clapping	386	157	71%
Bouncing on trampoline	534	315	63%
Baseball throw	293	256	53%
Climbing tree	438	390	53%
Cutting watermelon	27	723	4%
Squat	16	974	2%
Pull ups	59	704	8%

2) *Kinetics-Skeleton Subset (KSS)*: Kinetics-skeleton is a benchmark dataset used in HAR research that contains OpenPose-COCO extracted skeleton data of Kinetics-400 dataset. Due to the high computational cost associated with OpenPose extraction, this already extracted skeleton dataset was used in preliminary model building with Kinetics-400 and for the pre-training of ST-GCN model. Since majority of the Kinetics-400 data are shared data with Kinetics-600, implementations on child data were also done according to the following protocols using the shared-annotated data.

We introduce KSS protocols with the same naming convention (e.g., KSS-Full for full 8 subset implementation of kinetics-skeleton child data). KSS-Balanced protocol contains 110 samples per each of the 5 class.

¹Dataset and other resources: github.com/sankadivandya/KS-KSS-Dataset

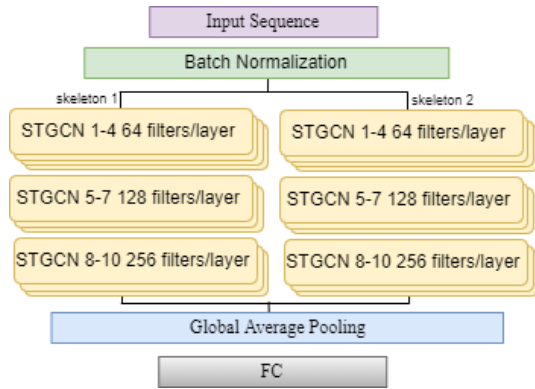


Fig. 1: ST-GCN Model Architecture

B. CAR on Skeleton Modality

TensorFlow based ST-GCN model implementation from our work [30] is taken as the base for this research. While the number of layers were kept same, number of filters was taken as a parameter that was determined in the first stage of the implementations. Feature vector in each node contains x, y pixel coordinates and the confidence value resulted in OpenPose. Since we are considering multi-person scenarios in our model, we run two ST-GCN models parallelly on each skeleton features but the convolution weights are shared among each model (Figure 1). The Global Average Pooling layer stacked on top of GCN layers combine the outputs for two skeletons in each video resulting in a single vector which is used as the input to the fully connected (FC) classifier.

Though tracking of skeletons is important in multi-person scenarios, as the actions considered are mostly single-person, we used a basic tracking mechanism in calculating the minimum Euclidean distance of skeletons in this paper. After normalizing the x, y feature values, centralization was done with changing the origin of coordinate system to the center of image. Camera movement was simulated as detailed in the ST-GCN paper [16] through affine transformations like rotation, translation and scaling. For videos with less than 10 seconds duration, zero padding was implemented. Zeros were assigned for missing joints and sorting of skeleton sequence in a multi-person scenario was also implemented.

1) *Learning Methods*: Model training was first done without any pre-trained data as detailed in Section IV-A2. To improve the model performance, it was pre-trained with kinetics-skeleton dataset which can be considered as a special case of propagation approach detailed in [30]. Furthermore, both feature extraction (FX) and fine-tuning (FT) methods [30] were implemented using kinetics-skeleton as the source dataset. Applying the best performance FT method [30], only the last GCN layer was kept trainable.

2) *Skeleton Structures*: OpenPose provides two pose estimation models called BODY_25 (i.e., OpenPose 2019 version) and COCO (i.e., OpenPose 2016 version). The BODY_25 is faster than COCO in extraction process and, its accuracy is also improved by 7% as detailed in [17]. COCO-skeleton (G_{coco}) contains 18 vertices(V_1) and BODY_25-




Fig. 2: Skeleton structures from OpenPose

skeleton (G_{body}) contain 25 vertices(V_2) and since $V_1 \subset V_2$ in $G_{coco} = (V_1, E_1)$ and $G_{body} = (V_2, E_2)$, G_{coco} skeleton structure could be used with the BODY_25 extracted data. Thus, skeleton extraction was done for the full kinetics600 subset using BODY_25 model.

Since kinetics-skeleton is created with COCO model and graph structure has to be same in all data, we used G_{coco} for KSS pre-trained model based implementations. Figure 2 compares the skeleton structures for a frame taken from a "climbing tree" class video and the black colour represent the skeleton in kinetics-skeleton dataset, blue for the COCO skeleton and red for the BODY_25 skeleton.

3) *Confidence Evaluations*: For each joint (J_n), OpenPose outputs a 3 element vector containing $[x_n, y_n, C_n]$ where C_n denotes confidence value and $C_n \in [0, 1]$. Though C_n can be used as a heuristic to analyse the effect of occlusion and truncation in videos misidentified skeletons and unusual poses can't be analysed. We calculate the average confidence of entire skeleton video for each person visible in the video. The calculation is further improved by removing skeleton-less frames and taking only the visible joints for the average. Extending these calculations to all sequences of a class, correlation of 'class-wise accuracy' and 'class-wise average confidence' variables were analysed using visualization, Pearson correlation and Spearman's correlation values.

C. CAR on RGB Modality

Action recognition in the original LRCN architecture used CNN and LSTM and was pre-trained on a subset of ImageNet dataset [11]. The LRCN approach used in this study was based on [31]. It was initially trained and tested with the UCF101 dataset [8] for a selected subset of 55 classes.

III. EXPERIMENTS

This section discusses the stages followed in conducting the quantitative analysis including pre-processing of the datasets, and the experimental settings.

A. Implementation Details

1) *ST-GCN*: The ST-GCN architecture was kept the same as in [16] in all implementations since early experiments with number of filters and layers didn't result in consistent improvements. A 'piece-wise constant decay' function was used as the learning rate scheduler where the learning rate was reduced with a rate of 0.1 after a given number of steps, while the initial learning rate and the other hyper-parameters were varied in each implementation. In vanilla implementations with KS dataset, base learning rate was 0.001 and for transfer learning (TFL) implementations it was 0.1. Though the learning rate was same for TFL implementations, it was varied with each vanilla implementation for KSS. Stochastic Gradient Descent (SGD) as the optimizer, categorical cross entropy as the loss function, and a batch size of 4 were used with number of epochs being either 30 or 50 in all implementations.

2) *LRCN*: The LRCN implementation applied in this research used the same parameters tested for UCF101 [31]. In the data preprocessing stage, sampling rate was 10, FPS of videos were 25, and the number of frames extracted was 15 for preprocessing of each video. The main configurable parameters of the model architecture were, the learning rate initially set to 0.0005, number of epochs 100, and the batch size was 16.

B. Evaluation methods

Top-1 accuracy as the main evaluation method together with confusion matrices were used to analyse the class-wise performance. Box-and-whisker plots were used to compare different learning methods in skeleton modality implementations and confidence interval calculations and visualization were done to compare different model performance.

IV. RESULTS AND DISCUSSION

Initial implementations were done to determine the primary configurations of the ST-GCN model. Then the secondary configuration (i.e., graph structure, pre-processing etc.) were selected for best performance and the resultant configurations were used throughout the rest of the implementations. Parallely, kinetics-skeleton implementation was done on ST-GCN to validate the model and to use as a pre-trained model.

A. Skeleton Modality Results

1) *Preliminary Implementations*: Since all the 8 interested activities are supposed to be done by a single person, performance was compared between one-person implementation (1-person Model) and two-people implementation (Standard Model) on KSS Full dataset. Implementation was further extended by discarding the confidence value (Section II-B3) from joint feature vectors (2D only Model). Since the standard model outperforms the 1-person model (Table II), 2 people per frame was chosen for number of people per frame configuration value. Though disregarding confidence feature values improves the accuracy (Table II), 2D+confidence (i.e., standard model) feature vector was chosen as the default configuration

TABLE II: Preliminary Implementation

Implementation	Top-1 Accuracy	Top-5 Accuracy
Standard Model	66.61	-
1-person Model	63.29	-
2D only Model	68.34	-
Kinetics-skeleton [original]	30.7	52.8
Kinetics-skeleton	21.16	41.7

since the improvement is marginal and we need to compare different key point extraction based models.

Full implementation of kinetics-skeleton was done but the data of the 8 classes were not used. Since that only account for 0.6% of kinetics-skeleton dataset, effect of this on model performance should be minimum. Original hyperparameters were used where it was possible but some (e.g., batch size) were not used due to computational limitations. This or the differences in pre-processing may have lead to the low performance in Table II.

For further comparison, implementations were done on the skeleton structure and between adult and child subsets as well. For these implementations, a balanced 4 class subsets were used with 98 samples per class. A modified graph structure instead of OpenPose-COCO graph structure where hip joints are connected to corresponding shoulder joints rather than the neck joint (Table III-Modified structure) was experimented which resulted in better performance. Thus, we attempted this configuration with the kinetics-skeleton pre-trained model which resulted in a sharp improvement of accuracy, both in child and adult dataset implementations (Table III-Modified pre-trained). Hence later implementations were attempted with kinetics-skeleton pre-trained implementations as well.

TABLE III: Secondary Implementation

Implementation	Child Dataset	Adult Dataset
OpenPose-COCO structure	63.26	46.93
Modified structure	61.22	56.12
Modified pre-trained	86.73	84.69
Random frame selection [$w = 150$]	55.10	43.87
Random skeleton movement	61.12	61.22
Combined approach [$w = 150$]	62.24	58.16
Sub-sampling approach	50.0	44.89

Experiments with final pre-processing stage were also done to select the best configuration. Authors of original ST-GCN paper [16] had implemented random frame selection process as well as random skeleton movement process, thus we first implemented them separately. Random frame selection was done with a window size (w) of 150 and gradually increased to 250, but the best results we achieved ($w = 150$) were below the performance of "Modified structure" approach (Table III). As the best performance was achieved when both processes were combined (Table III-Combined approach), it was considered as the default configuration. In addition, a new sub-sampling approach with frame dropping was also implemented, but due to low performance it was not added to the final pre-process stage.

For KS based implementations, comparisons were done between different skeleton structures (Table IV). Feet related

joints were removed from G_{body} resulting in G_{body}^* skeleton structure with 19 vertices (V_3). Considering the overall performance, G_{body} was used for the vanilla implementations while G_{coco} was used with pre-trained model implementations.

TABLE IV: Skeleton Structure Selection

Skeleton structure	KS-balanced	KS-Full
$G_{body}, V_2 = 25$	69%	75%
$G_{body}^*, V_3 = 19$	67%	74%
$G_{coco}, V_1 = 18$	66.45%	74%

2) *Vanilla implementation on Child data:* Implementations were done on both KS and KSS protocols and the models were not pre-trained on any other dataset. Results in Table V suggest a general improvement in accuracy when moving from a KSS protocol to a KS protocol. This improvement could be due to the higher number of data samples per class in KS or due to the improvement of graph structure as demonstrated by the results in Table IV. When comparing the class-wise accuracy between KS/KSS-Balanced protocols, performance of each class has increased, yet there is no relative improvement between classes. ‘Clapping’ class perform the best while ‘baseball throw’ perform the worst. Result also suggest there is no strong connection between confidence value and accuracy given that ‘climbing tree’ class performs second best even though the average confidence value is the lowest and ‘hopscotch’ performs second worst even though the average confidence is the highest.

TABLE V: Vanilla Implementation Results

From Scratch	Full	Large	Balanced	Small-C	Small-A
KS	75.29	77.83	69.32	69.23	92.34
KSS	60.68	64.88	59.43	86.95	-

3) *Kinetics-Skeleton Implementation:* Class-wise evaluation of the kinetics-skeleton implementation introduced in Section IV-A1 was done and the results are given in the Table VI for the 8-class subset. Class index refers to the class index of the Table I. ‘Position’ refers to the place each class take when all 400 classes are ordered in descending order in terms of class-wise accuracy. Confidence value is calculated per person in video and since there are only 2 people maximum, both of these values are given in this row.

Higher accuracy and position attained by classes with indices 7, 6, 0, and 2 (Table I) can be explained as a result of motion-oriented nature of those actions. Considering the distribution of all Kinetics-skeleton 400 classes in accuracy vs confidence (Figure 3), all four classes are above average. Relatively bad performance of other 4 classes is difficult to attribute to a single cause. Considering 4 and 5 classes, it may be due to truncation/occlusion present in the videos as evidenced by the confidence values but same reasoning is not true for the low performance of 1 and 3 classes.

Visualized distribution of all 400 classes in Figure 3 implies there is no strong correlation but if the average confidence values is close to zero, then there is a higher chance of resulting in a low accuracy. Quantitative analysis resulted

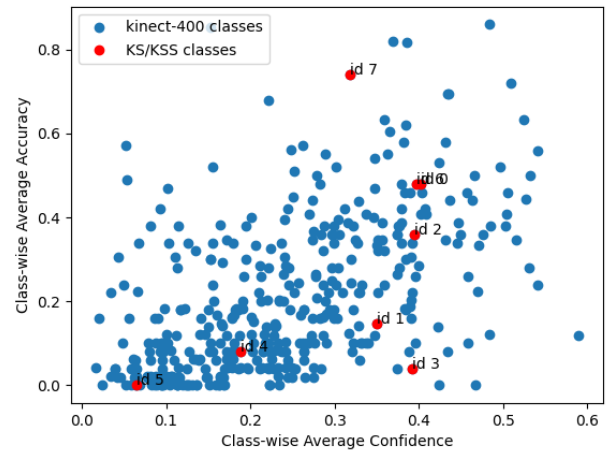


Fig. 3: Accuracy vs Confidence Comparison

in 0.533 for Pearson correlation indicating only a moderate positive relationship and 0.564 for Spearman’s correlation. Analysis on other implementation also resulted in similar results and conclusions.

TABLE VI: kinetics-skeleton results for 8 classes

Class Index	0	1	2	3	4	5	6	7
Accuracy	48%	14.58%	36%	4%	8%	0.00%	48%	74%
Position	32(2)	210(5)	79(4)	321(7)	276(6)	373(8)	34(3)	5(1)
Confidence	.40/.12	.35/.18	.39/.11	.39/.16	.19/.03	.06/.01	.40/.12	.32/.05

4) *Transfer Learning Implementations:* Initial implementations with propagation approach resulted in considerable accuracy improvement for both KS and KSS implementations. This can be attributed to the Kinetics-skeleton dataset size and diversity. FX and FT approaches using Kinetics-skeleton as the source dataset resulted in marginal differences compared to propagation approach. But in our previous work [30], both FX and FT outperformed propagation approach and this could be either due to increased diversity in source dataset in this implementation compared to the NTU dataset used as the source dataset in [30] or simply due to the presence of negative TFL in these implementations. Further experiments are needed to validate these claims.

To analyse these approaches in detail, KSS-Full protocol class-wise results for each approach was visualized using a box-and-whisker plot in Figure 5. Resultant probability value from ST-GCN model’s final softmax-activation based dense layer was used as the sample probability to develop the probability distribution of each class. Transfer Learning approaches class-wise gain over vanilla approach is evident through the median values of box plot. Based on the median values, propagation approach performs better than other TFL methods even though the difference is small. When comparing the distribution variance of each approach, propagation approach has a considerably low variance in ‘pull ups’ class and the large five classes with the exception of ‘baseball throw’ class. Since all these distributions can be interpreted as combinations of correctly classified sample distribution and misclassified sample distribution, the high variance of

TABLE VII: Transfer Learning Results

	Implementation	Full	Balanced	Large	Small-C
KS	Propagation	84.3	83.38	86.03	76.92
	Propagation	81.26	87.68	87.92	82.60
KSS	Fine-Tuning	80.47	89.85	86.51	82.60
	Feature Extraction	79.15	89.13	87.92	86.95

‘throw baseball’ can be explained as a result of considerable contribution of misclassified sample distribution. While the class-wise accuracy improvement of 5 and 6 classes are negligible, with propagation approach, upper quartile has increased considerably, implying there is a room for improvement if more data is present.




Fig. 4: Probability Distribution of Skeleton Modality Implementations

B. RGB Modality Results

RGB implementations were carried out on the same dataset protocols as given in Section II-A. The performance comparison between different datasets shows a good accuracy in RGB implementation except for KS-Small-C dataset due to low number of samples(Table VIII). When compared to the other datasets, the balanced dataset achieves the best accuracy for both KSS and KS protocol implementations. Furthermore, the use of pre-trained ResNet-152 in LRCN enhanced the model’s training efficiency and prevented overfitting for lower number of samples.

TABLE VIII: Comparison of RGB modality and Skeleton modality

Dataset	Modality	Full	Balanced	Large	Small-C
KS	RGB	86.62	88.64	87.02	73.07
	Skeleton	84.3	83.38	86.03	76.92
KSS	RGB	82.57	86.23	79.72	78.26
	Skeleton	81.26	89.85	87.92	86.95

C. Model Comparison

Since LRCN model is pre-trained on the ImageNet dataset, comparison of ST-GCN vanilla implementation with LRCN is not sensible. Instead we compare the pre-trained ST-GCN model based implementations (i.e., propagation approach) with the pre-trained LRCN implementation. Between the KS/KSS protocol based implementations, the differences are marginal as detailed in Table VIII. While the LRCN performs better than

TABLE IX: Comparison of Models

Model	Kinect Camera	OpenPose Model	RGB Model
Accuracy	70.56	84.3	86.62

ST-GCN in KS-Balanced protocol, the opposite is true regarding the KS-Small protocol. Comparing the KSS protocols, ST-GCN performs better than LRCN in both KSS-Balanced and KSS-Small-C protocols. Since the differences are marginal, we argue that performance of both models are comparable.

A class-wise sample probability comparison was done with confidence interval (CI) of 95% for KSS-Balanced protocol. Result in Figure 5 suggest that LRCN performs better in ‘hopscotch’ and ‘climbing tree’ classes but ST-GCN performs better in all others. Thus, even though in terms of Top-1 accuracy, ST-GCN performs better, these results also show that overall performance is similar. In some instances upper limit of CI goes beyond 1 as a result of Gaussian assumption in CI calculation.

In order to compare 2D Skeleton and RGB modality results with similar classes of 3D data, for KS/KSS-Full protocol, 8 similar class implementations using the CWBG dataset were done with the ST-GCN model used in [30] and the best result was achieved with FT approach. The 3D result is shown as ‘Kinect-Camera’ in Table IX along with best results we achieved in this work.




Fig. 5: Performance of Models with Confidence Interval

V. CONCLUSION AND FUTURE WORKS

Inspired by the superior performance of skeleton based models in HAR research, we carry out an in-depth analysis of CAR on pose-based GCN. This analysis is further extended with a comparison of RGB modality and Skeleton modality in CAR. Though the average class-wise confidence values show that unconstrained nature of videos severely limit the pose estimation process, in this research we show that ST-GCN model is able to achieve comparable performance to LRCN model.

Based on our results on KS/KSS-Full protocol, LRCN performance over ST-GCN suggest a limitation of GCN when

truncation and occlusion are present in videos. ST-GCN performance over LRCN on KS/KSS-Small-C protocol suggests a better discrimination ability of GCN when presented with small number of classes. On KS/KSS-Balanced protocol both models perform equally well suggesting similar potential of LRCN and ST-GCN model architectures.

While the average confidence value intuitively quantify the ‘visibility’ of the skeleton sequence, Pearson correlation coefficient value of 0.53 between class-wise accuracy and average confidence value suggest that in addition to the average confidence value, presence of other latent variables affect skeleton modality performance.

Despite the pose differences in adults and children, these results suggest that when the actions are motion oriented, skeleton modality can perform on par with RGB modality, thus opening future research directions in optimal skeleton-graph creation, pose estimation improvements in complex scenarios and developing ensemble of multiple modalities.

ACKNOWLEDGMENT

This research was supported by the Accelerating Higher Education Expansion and Development (AHEAD) Operation of the Ministry of Higher Education of Sri Lanka funded by the World Bank (<https://ahead.lk/result-area-3/>).

REFERENCES

- [1] E. E. Aksoy, Y. Zhou, M. Wächter, and T. Asfour, “Enriched manipulation action semantics for robot execution of time constrained tasks,” in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, 2016, pp. 109–116.
- [2] D. Singh and C. Krishna Mohan, “Graph formulation of video activities for abnormal activity recognition,” *Pattern Recognition*, vol. 65, pp. 265–272, 2017.
- [3] Y. Gao, X. Xiang, N. Xiong, B. Huang, H. J. Lee, R. Alrifai, X. Jiang, and Z. Fang, “Human action monitoring for healthcare based on deep learning,” *IEEE Access*, vol. 6, pp. 52 277–52 285, 2018.
- [4] N. Crasto, P. Weinzaepfel, K. Alahari, and C. Schmid, “Mars: Motion-augmented rgb stream for action recognition,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 7874–7883.
- [5] J. Goto, T. Kidokoro, T. Ogura, and S. Suzuki, “Activity recognition system for watching over infant children,” in *2013 IEEE RO-MAN*, 2013, pp. 473–477.
- [6] T. L. Westeyn, G. D. Abowd, T. E. Starner, J. M. Johnson, P. W. Presti, and K. A. Weaver, “Monitoring children’s developmental progress using augmented toys and activity recognition,” *Personal Ubiquitous Comput.*, vol. 16, no. 2, p. 169–191, feb 2012.
- [7] G. Sciortino, G. M. Farinella, S. Battiato, M. Leo, and C. Distanto, “On the Estimation of Children’s Poses,” in *Image Analysis and Processing - ICIAP 2017*, S. Battiato, G. Gallo, R. Schettini, and F. Stanco, Eds. Cham: Springer International Publishing, 2017, pp. 410–421.
- [8] K. Soomro, A. R. Zamir, and M. Shah, “Ucf101: A dataset of 101 human actions classes from videos in the wild,” *arXiv preprint arXiv:1212.0402*, 2012.
- [9] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev, M. Suleyman, and A. Zisserman, “The kinetics human action video dataset,” *CoRR*, vol. abs/1705.06950, 2017.
- [10] J. Carreira, E. Noland, A. Banki-Horvath, C. Hillier, and A. Zisserman, “A short note about kinetics-600,” *arXiv preprint arXiv:1808.01340*, 2018.
- [11] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell, “Long-term recurrent convolutional networks for visual recognition and description,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 2625–2634.
- [12] V. Veeriah, N. Zhuang, and G.-J. Qi, “Differential recurrent neural networks for action recognition,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4041–4049.
- [13] B. Fernando, E. Gavves, J. Oramas, A. Ghodrati, and T. Tuytelaars, “Modeling video evolution for action recognition,” *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5378–5387, 2015.
- [14] 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.
- [15] L. Shi, Y. Zhang, J. Cheng, and H. Lu, “Two-stream adaptive graph convolutional networks for skeleton-based action recognition,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 12 026–12 035.
- [16] S. Yan, Y. Xiong, and D. Lin, “Spatial temporal graph convolutional networks for skeleton-based action recognition,” in *Thirty-second AAAI conference on artificial intelligence*, 2018.
- [17] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, “Realtime multi-person 2d pose estimation using part affinity fields,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 7291–7299.
- [18] V. Bazarevsky, I. Grishchenko, K. Raveendran, T. Zhu, F. Zhang, and M. Grundmann, “Blazepose: On-device real-time body pose tracking,” *arXiv preprint arXiv:2006.10204*, 2020.
- [19] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, “NTU RGB+D: A large scale dataset for 3D human activity analysis,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1010–1019.
- [20] T. S. Kim and A. Reiter, “Interpretable 3d human action analysis with temporal convolutional networks,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017, pp. 1623–1631.
- [21] Y. Zhang, Y. Tian, P. Wu, and D. Chen, “Application of skeleton data and long short-term memory in action recognition of children with autism spectrum disorder,” *Sensors*, vol. 21, no. 2, 2021.
- [22] S. Suzuki, Y. Amemiya, and M. Sato, “Enhancement of gross-motor action recognition for children by cnn with openpose,” in *IECON 2019-45th Annual Conference of the IEEE Industrial Electronics Society*, vol. 1. IEEE, 2019, pp. 5382–5387.
- [23] Y. Amemiya, S. Suzuki, and M. Sato, “Enhancement of child gross-motor action recognition by motional time-series images conversion,” in *2020 IEEE/SICE International Symposium on System Integration (SII)*. IEEE, 2020, pp. 225–230.
- [24] R.-D. Vatavu, “The dissimilarity-consensus approach to agreement analysis in gesture elicitation studies,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1–13. [Online]. Available: <https://doi.org/10.1145/3290605.3300454>
- [25] Y. Dong, A. Aristidou, A. Shamir, M. Mahler, and E. Jain, “Kinder-Gator 2.0, Optical motion capture, Dataset, MIG2020,” Oct. 2020.
- [26] S. S. Rajagopalan, A. Dhall, and R. Goecke, “Self-stimulatory behaviours in the wild for autism diagnosis,” in *2013 IEEE International Conference on Computer Vision Workshops*, 2013, pp. 755–761.
- [27] A. Aloba, G. Flores, J. Woodward, A. Shaw, A. Castonguay, I. Cuba, Y. Dong, E. Jain, and L. Anthony, “Kinder-Gator: The UF Kinect Database of Child and Adult Motion,” in *EG 2018 - Short Papers*, O. Diamanti and A. Vaxman, Eds. The Eurographics Association, 2018.
- [28] F. Baradel, C. Wolf, J. Mille, and G. W. Taylor, “Glimpse clouds: Human activity recognition from unstructured feature points,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 469–478.
- [29] H. Duan, Y. Zhao, K. Chen, D. Lin, and B. Dai, “Revisiting skeleton-based action recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 2969–2978.
- [30] S. Mohottala, P. Samarasinghe, D. Kasthurirathna, and C. Abhayaratne, “Graph neural network based child activity recognition,” in *International conference on Industrial Technology*. IEEE, In Press.
- [31] doronharitan, “Action Recognition in Videos Using LRCN,” 01 2020. [Online]. Available: https://github.com/doronharitan/human_activity_recognition_LRCN