This WACV 2020 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Spatio-Temporal Pyramid Graph Convolutions for Human Action Recognition and Postural Assessment

Behnoosh Parsa* University of Washington behnoosh@uw.edu Athma Narayanan^{*} Behzad Dariush Honda Research Institute {anarayanan, bdariush}@honda-ri.com

Abstract

Recognition of human actions and associated interactions with objects and the environment is an important problem in computer vision due to its potential applications in a variety of domains. Recently, graph convolutional networks that extract features from the skeleton have demonstrated promising performance. In this paper, we propose a novel Spatio-Temporal Pyramid Graph Convolutional Network (ST-PGN) for online action recognition for ergonomics risk assessment that enables the use of features from all levels of the skeleton feature hierarchy. The proposed algorithm outperforms state-of-art action recognition algorithms tested on two public benchmark datasets typically used for postural assessment (TUM and UW-IOM). We also introduce a pipeline to enhance postural assessment methods with online action recognition techniques. Finally, the proposed algorithm is integrated with a traditional ergonomics risk index (REBA) to demonstrate the potential value for assessment of musculoskeletal disorders in occupational safety.

1. Introduction

Human action recognition has been a widely studied research topic in computer vision for several decades. The task is to infer the human action and activity from still images or video frames. Solutions to this important and challenging problem have traditionally been applied to domains such as surveillance, entertainment, robotics, video retrieval, and intelligent driving assistance systems [38, 34, 64]. Recently, there are emerging applications that involve assessment of human performance for virtual fitness, health monitoring, training, and ergonomics risk assessment for occupational safety [35, 10, 40]. These applications have unique requirements that may involve simultaneous association of time varying pose with action and object interaction, and relating such information for computational mod-



Figure 1: Our model (ST-PGN) takes a sequence of skeleton input produced by a pose extraction unit (like LCR-Net [42]) and does early action recognition. The skeleton sequence along with the activity labels go to the REBA computation unit to assess the ergonomics risk while testing.

eling and prediction of various biomechanical indicators. Vision only systems are non-invasive and less expensive alternatives to study these problems as opposed to expensive drift prone motion capture systems and wearable sensors [33, 5]. Depending on the application, human action recognition can be formulated in an online or off-line setting. In most applications, processing is performed off-line, making use of the entire video sequence without strict limitations on computational resources. In such cases, the typical assumption is that the start and end points of the action is known [9, 32] and the training video is pre-segmented into various action classes. Recent advances in hardware and GPU performance has led to the emergence of many online applications, where the requirement is to process video streams in real-time and without a priori knowledge of the transitions between actions [28, 29, 50]. Generalization of action recognition algorithms is a challenging and unsolved problem. Ideally, the method should generalize to various environments and deal with cluttered backgrounds, occlusions, and viewpoint variations. While end to end video to action

^{*}Co-first authors who contributed equally to this work.

This work was done while Behnoosh Parsa was an intern at Honda Research Institute USA.



Figure 2: The Feature Pyramid Convolutional Graph Network pipeline.

classification have shown great promise, generalization is achieved through domain adaptation [11, 2] or using intermediate skeletal representations that are robust to these variations [60, 41]. In particular, skeleton-based features appear to produce favorable results since human pose is typically a consistent representation of action across people and context. Among them, recent work based on graph convolutional networks that extract meaningful features from the skeleton have achieved good performance [60, 24].

The work in this paper is inspired by emerging applications involving human performance assessment in various domains including health, fitness, rehabilitation, and occupational safety. In particular, we consider specific challenges for real-time ergonomics risk assessment in complex environments such as manufacturing assembly. The requirements include correlation of action with the time varying posture and associated ergonomics and biomechanical risk. The ultimate goal is to produce reliable estimates of pose, action, and associated ergonomics indicators in order to identify the risk of musculoskeletal disorders associated with acute and repetitive tasks.

Contributions of this work. We propose a novel realtime Spatio-Temporal Pyramid Graph Convolutional Network (ST-PGN) for action recognition that enables the use of features from all levels of the skeleton feature hierarchy (Figure 2). ST-PGN, designed with a feature pyramid architecture enables the model to capture the correlation between body parts, rather than hand-coding body-part relations. We test the performance of the model on two public benchmark datasets typically used for postural assessment (TUM and UW-IOM) as well as Kinetics and NTU-RGBD datasets. We show that the algorithm is also able to learn the transitions between actions and is suitable for real-time applications. As compared to the state-of-the-art algorithms such as ST-GCN [60], our model has fewer graph convolution kernels without sacrificing performance. Finally, we enhance the pipeline with postural assessment methods (REBA [13]) that use the online action recognition output of our model to produce ergonomics risk estimates. We propose, this combined action-risk architectural design as a first step towards automated assessment of musculoskeletal disorders in occupational safety.

2. Related Work

Given the recent advances in obtaining accurate pose through depth sensing or vision based pose estimation algorithms, skeleton based action recognition methods have become vital for achieving generalization across a variety of environments [42]. Skeleton based methods also offer opportunity to study down stream applications that involve human performance analysis and require postural assessment. We summarize work related to the proposed ST-PGN algorithm for association of action, posture, and ergonomics risk. To the best of our knowledge this is the first work that combines the three separately studied problems jointly and in an online fashion.

2.1. Action Classification

Video action segmentation tasks focus on localizing action labels in untrimmed videos or classifying entire video clips with one label [46, 52, 63, 55, 4, 54, 60]. For further literature on action recognition please refer to [57]. Usually large datasets are used for evaluating algorithms developed for these tasks [51, 20, 16, 4]. Variations of temporal convolutions outperform conventional recurrent networks in these tasks due to its ability of aggregating long-term motion changes [1, 21]. Extending these models to an ongoing, partially observed, and multi-action sequence, such as online ergonomics risk assessment task, is unclear. Hence, we focus on models and datasets (TUM [53] and UW-IOM [35]) that translate to ergonomics risk assessment settings that contain repetitive actions that are closely tied to activities such as manufacturing assembly.

With advances in reliable pose estimation models [62, 42, 3], skeleton only action recognition has gained popularity [61, 60, 22, 48]. Those methods have shown to be robust to variations in illumination and scene, and are typically context agnostic. One limitation of previous methods is that they do not use the necessary features from scene context or object handling which give more meaning to the actions (e.g. walking on crosswalk means crossing verses walking indoors, lifting box vs lifting rod). Using scene only cues limits models from capturing complex pose dynamics and relative pose structure changes (e.g. hand moving in relation to torso means reaching for object). In this regard [60] is, to our knowledge, the first method to operate on a local pose structure graph. Our work addresses the subproblem of online action classification for ergonomics risk assessment by exploiting hierarchical spatio-temporal structure. To avoid comparison to plethora of action recognition work, we compare our models to spatio-temporal models that use GCN.

Most similar to our work are [28, 29, 50], which address online early action recognition for indoor datasets such as [25, 27]. However, the focus of those works is on modelling the temporal evolution of poses and early prediction of future actions. Rather than predicting future pose streams, our aim is to instead classify incoming pose streams. It is imperative to capture local label transitions (*reaching* to *pickup*) by exploiting subtle pose cues and temporal sequence understating. Moreover, as evidenced by our ST-GCN [60] experiments, offline models do not translate well to online settings. Hence, we design a hierarchical architecture that can do these tasks jointly in an online manner.

2.2. Graph Convolution Networks

Graph convolution network (GCN) is a powerful method for processing non-Euclidean spaces [58]. Since the skeleton structure is inherently represented as a graph with nodes and connections, GCN is increasingly being used for analyzing human motion for different applications. Spatiotemporal graph convolutions add another dimension to GCN by applying convolutions over spatial domain, and temporal convolutions (TCN) over the time domain in a sequential manner. Most related work in skeleton based action recognition include [60, 17, 24, 48]. The first three papers focus on graph convolution on temporal skeleton sequences. However, they do not model the hierarchical parts structure in graphs.

Recently, Kim et al. [17] introduced a two-stream method for human action recognition. They used a human pose stream based on ST-GCN and an object-related pose stream which is achieved by training an object detector on the set of objects of their interest. Similarly, our work attempt to fuse the object/context features along with pose dynamics. However, we focus on enhancing the skeleton features and treat objects as features from VGG16. We propose an alternative strategy for fusion inspired by GRU. The focus is to avoid confusion between objects handled in the labels (pose configuration for *rod-pickup* and *box-pickup* look similar).

2.3. Ergonomics Risk Assessment

Work-related musculoskeletal disorders (MSDs) are costly, affect all age groups, and are common in many occupations. MSDs are a major contributing factor to disability, loss of independence, and early retirement. Therefore, many studies analyze the ergonomics risk for workers, particularly in manufacturing assembly [37, 49, 6, 43, 39, 30]. Rapid Entire Body Assessment (REBA) [14] and the European Assembly Worksheet (EAWS) [44] are two common ergonomics risk measures used in the industry. REBA is a tabular method created by experts in ergonomics by evaluating over 600 postural examples. REBA is a less qualitative measure for ergonomics risk assessment which takes the human joint angles and computes a risk score. EAWS, however, is focused at the type of activity that is done during an assembly task. Both metrics are traditionally determined visually, by an expert observing the action.

One line of research is focused on using body-mounted motion sensors for automation of ergonomics risk assessment [23, 30]. Recently, [31] introduces a dataset which is very useful for studying ergonomics for collaborative robotics application. Another track focuses on using only a camera sensor for evaluating the safety of an activity. For instance, in [45] an RGB-D camera is used to find a safe range for arm movements and give feedback on the subjects' performance during welding. In [18] a camera is used to monitor and adjust the ergonomics risks of working with power tools in real-time. Moreover, in [35], an offline method is introduced to segment a video into semantically meaningful actions and report an ergonomics risk level for each action.

Improved ergonomics risk assessment can be attained by considering not only the posture, but also the action and object interaction. In this work, we compute REBA framewise and use the recognition predictions to adjust the scores. Our activity recognition predicts the postures and actions, and identifies object interactions and the height at which the activity is being performed, which are important for REBA calculation. For additional details on REBA computation, refer to [14].

3. Spatio-Temporal Feature Pyramid Graph Convolution

In this work we introduce Saptio-Temporal Pyramid Graph Convolutional Network (ST-PGN). ST-PGN models the spatio temporal features of the skeletal structure using combinations of Pyramidal GCNs (PGNs) and Long-Short-Term-Memory Units (LSTMs). PGN is a novel way to process non-Euclidean skeletal data in a hierarchical form. Each feature representation in PGN hierarchy is used as an input to an LSTM unit to learn the temporal aspect of the input sequence (shown in Figure 2 and described in Section 3.4).

3.1. Graph Convolutional Network

Graph convolutional networks (GCN) [65] learn the layer-wise propagation operation that can be applied on structured data represented by a graph. To briefly introduce how GCNs work, assume we have an undirected graph with N nodes, a set of edges between nodes, an adjacency matrix $\mathbf{A} \in \mathbf{R}^{N \times N}$, and a degree matrix $\mathbf{D}_{ii} = \sum_{j} \mathbf{A}_{ij}$. If $\mathbf{x} \in \mathbf{R}^{F \times N}$ represents the feature matrix of the graph $(\mathbf{x}_i \in \mathbf{R}^F)$ is the feature vector of node *i* with size *F*), a linear formulation of graph convolution is,

$$\mathbf{f} = \hat{\mathbf{D}}^{-\frac{1}{2}} \hat{\mathbf{A}} \hat{\mathbf{D}}^{-\frac{1}{2}} \mathbf{x}_i^{\mathsf{T}} \mathbf{W}, \tag{1}$$

where $\hat{\mathbf{A}} = \mathbf{A} + \mathbf{I}$, \mathbf{I} is the identity matrix and $\mathbf{W} \in \mathbf{R}^{F \times C}$ is the weight matrix. So, if the input to a GCN layer is $F \times N$ the output feature \mathbf{f} is $N \times C$, where C is the chosen output size. As with any other convolution layer we can have a stack of GCNs each followed by a nonlinear function (such as ReLU) [19].

In this work, we are following the spatial configuration partitioning introduced in ST-GCN [60], therefore, $\hat{\mathbf{A}} = \sum_{a} \mathbf{A}_{a}$ and Equation 1 is written in a summation form.

$$\mathbf{f} = \sum_{a} \hat{\mathbf{D}}_{a}^{-\frac{1}{2}} \mathbf{A}_{a} \hat{\mathbf{D}}_{a}^{-\frac{1}{2}} \mathbf{x}^{\top} \mathbf{W}_{a}, \qquad (2)$$

Equation 2 is represented for k^{th} level of the pyramidal hierarchy in line 4 of Algorithm 1. We hypothesize that a hierarchical graph convolution that operates on human joints, body parts and global structure would enrich the input representation.

3.2. Pyramidal Graph Architecture

Pyramidal Graph Convolutional Network (PGN) is a hierarchical GCN that produces different spatial features with semantic meaning at different levels. The input to the PGN is the skeleton with N joints represented by a tensor (X) of dimension $F \times N \times T$, where T indicates time. Each GCN aggregates features along the spatial dimension using a specific adjacency matrix A_k using Equation 2. Our PGN has three graph levels $(\hat{A}_1, \hat{A}_2, \hat{A}_3)$. The initial GCN works on the skeleton with \hat{A}_1 , which is constructed based on the skeleton connections and accompanied with an edgeimportance matrix ¹. The subsequent graph levels represent the body parts and global structure respectively. Since the correlation between the nodes for higher level graphs is unknown, $\hat{\mathbf{A}}_2$ and $\hat{\mathbf{A}}_3$ represent fully connected graphs and we let the edge-importance learn the correlations. Thus our model has a hierarchy of graphs with the base as the input skeleton and the top level a graph with three nodes representing arms, legs and middle part of the body. We refer to this hierarchical graph structure as a pyramidal graph architecture because it is large at the base and becomes smaller as we move to the top levels.

Symbol	Legend
N	Number of 3D Skeleton joints, (x, y, z) tuples
T	Time history of 80 samples
\mathbf{c}_k	Graph convolution(GCN) at each hierarchy k
$\hat{\mathbf{A}}_k$	Adjacency matrix at each hierarchy k
\mathbf{X}_0	Input Skeleton feature to first GCN ($3 \times N \times T$)
\mathbf{g}_k	Group Average Pool at each hierarchy k
\mathbf{J}_k	Pooling kernel at each hierarchy k
\mathbf{f}_k	Output of each GCN ($3 \times N \times F$)
\mathbf{w}_k	1×1 convolution operation
\mathbf{u}_k	Upsample and Add
\mathbf{p}_k	Output of 1×1 convolution
\mathbf{z}_k	Final features sent to LSTMs

Table 1: Description of the symbols used in Algorithms

3.3. Group Average Pool

A Group Average Pool (GAP) layer average-pools the features in a selected group of nodes or joints using a specific kernel (\mathbf{J}_k) for each level (line 5 of Algorithm 1). The resulting graph has nodes that represent a higher level body part (as shown in Figure 2). Therefore, every layer of the pyramid has a semantic meaning, from low to high level. In the bottom left corner of Figure 2, we show how the groups are defined in TUM and UW-IOM datasets.

More specifically, feature masking is inspired by [7] which is generally used in foreground background separation. Here, kernels are pre-determined matrices with ones or zeros. These kernels are element-wise multiplied by the features to group only certain body parts one at a time. For example, the kernel has ones in the particular rows corresponding to those joints representing left arm (7, 9, 11) and zero everywhere else. Hence, the masked features (features multiplied by the mask) all belong to the left arm.

¹Edge-importance is a learnable mask on every layer of the spatiotemporal graph convolution. It has the same dimension as the adjacency matrix and learns to scale the contribution of a nodes feature to its neighboring nodes based on the learned importance weight of each spatial graph edge.

These features are average pooled as they belong to the same group. Multiple such combinations are used to group the joints into different parts. Similarly, parts are combined into global structure using another set of kernels. Such successive GCN-GAP combinations allows us to model the entire local and global motions jointly. We refer to this as the feature update rule (Algorithm 1), and later in Section 3.4, it is referred to as a bottom-up pathway.

3.4. Feature Pyramid Graph Convolutional Network

Feature pyramids have been an important component of object recognition algorithms [26, 47, 12]. The advantage of using pyramids is that it produces a multi-scale feature representation in which all feature levels are semantically strong. Especially in skeleton-based action recognition the correlation of body-parts can be very informative in recognizing actions. However, a pre-defined graph might not be sufficient to represent every sample. For example, in ST-GCN graph, there is no connection between hand and head, which is important in actions such as eating. Therefore, here we are generalizing the feature pyramid network to a GCN pyramidal feature hierarchy, and we believe that learning the correlations at different levels of the hierarchy enhances the performance of our model. Here feature pyramids are still valid in skeleton structure as global motion is a combination of local motion of parts and part motion is a combination of local motion of joints. Hence our feature pyramids aggregate joints, parts and global features jointly [59]. The feature pyramid networks consist of two pathways, a

Algorithm 1 Feature Update Rule								
1: $\mathbf{X}_0 \leftarrow \mathbf{X}$	▷ input skeleton distributed over time							
2: $k \leftarrow 1$	⊳ iterator							
3: while $k \leq 3$ do								
4: $\mathbf{f}_k = \mathbf{c}_k(\mathbf{X}_{k-1}; \hat{\mathbf{A}}_k)$	▷ GCN operation							
5: $\mathbf{X}_k = \begin{cases} \boldsymbol{g}_k(\mathbf{f}_k; \mathbf{J}_k) \\ None \end{cases}$	if $k < 3$, otherwise. \triangleright GAP operation							
6: k=k+1								
\triangleright The $\mathbf{f}_k \forall k \in (1, 2, 3)$ are used as input features for the feature pyramid operations in Algorithm 2.								

Alg	Algorithm 2 Pyramid Update Rule							
1:	$k \leftarrow 1$	⊳ iterator						
2:	while $k \leq 3$ do							
3:	$\mathbf{p}_k = \begin{cases} \boldsymbol{w}_k \otimes \mathbf{f}_k & \text{if } k < 3, \\ \mathbf{f}_k & \text{otherwise.} \end{cases}$	$\triangleright 1 \times 1$ convolution						
4:	$oldsymbol{z}_k = egin{cases} oldsymbol{p}_k \oplus oldsymbol{u}_k(oldsymbol{p}_{k-1}) & ext{if } k < 3, \ oldsymbol{p}_k & ext{otherwise.} \end{cases}$							
5:	× ×	⊳ Upsample & Add						
6:	k=k+1							
	\triangleright The Following $z_k \forall k \in (1, 2, 3)$ are used temporal modelling using three separate LS	as input features for the TMs.						

bottom-up and a top-down pathway. The bottom-up path-

way is the feed-forward computation of the backbone GCN, which computes a feature hierarchy consisting of feature maps at different scales. The **top-down pathway** produces higher resolution features by up-sampling spatially larger, but semantically stronger, feature maps from higher pyramid levels. The top-down path is enhanced by the features produced in the bottom-up pathway through lateral connections. The features from the bottom-up pathway undergo a 1×1 conv layer to reduce channel dimensions and then are merged into the top-down pathway features by elementwise addition. The purple connections in Figure 2 shows this process, and it is described as the pyramid update rule.

3.5. Spatio-Temporal Modelling

Now we briefly summarize ST-PGN steps that are described in Algorithm 1-2 and Figure 2, and also describe major differences with respect to ST-GCN. The input skeleton (\mathbf{X}_0) goes through three levels of GCN and GAP, and the output of each level (\mathbf{f}_k) is aggregated with the upsample features through lateral connection and forms the final features (\mathbf{z}_k) . Each pyramidal feature is passed through separate LSTMs to create three frame-wise activity predictions. As an ablation study we either 1) average these three predictions and compute one loss or 2) compute three losses separately and average the predictions while testing. The latter gives us better performance. As a comparison, in ST-GCN, the input goes through a sequence of multiple GCN and TCN units so that the final feature embodies spatial and temporal properties of the input. A final feature that summarizes spatial and temporal properties is the key for video clip classification. However, when we need to recognize activities frame-wise, that strategy fails as it is shown in Section 5.2. Therefore, we extract the spatial features through PGN and send these features to individual LSTM units so that the temporal aspect is learned at different spatially semantic layers.

3.6. Optional Fusion Unit

To study the benefit of image features, we also perform experiments with image features concatenated along with skeleton pose features. We hope to avoid confusion in situations with object handling. Hence we extract VGG16 features from a crop image region around the human and fuse them with the final skeleton feature pyramid. Our fusion unit is inspired by GRU [8], that learns to weight the features before LSTM. We freeze the weights of the pre-trained network and only train the fusion unit along with the final LSTM layer. While the benefit of the image features are very minimal, for completeness, we will describe the fusion unit below.

At time t, let the image features and the final feature pyramid layer features be denoted by i_t and z_{1t} , respectively. Since the dimensions of these features do not match, we apply linear weights (U_i, U_z) to transform them into the same dimension and arrive at the transformed image and skeleton features I_t and S(t) as shown in Equation 3. The terms W_i and W_z are learnt weights that are used to learn a gauging value (p_t) between the two features similar to the GRU. The weight p_t is squished to take on values $\in [0, 1]$ using a sigmoid operation. Finally this weights are multiplied to the incoming features.

$$I_t = \operatorname{relu}(\mathbf{U}_i * i_t), S_t = \operatorname{relu}(\mathbf{U}_z * z_{1t})$$
(3)

$$p_t = \sigma(\mathbf{W}_i * I_t + \mathbf{W}_s * S_t), \tag{4}$$

$$\mathbf{O}_t = p_t I_t + (1 - p_t) S_t \tag{5}$$

Where, O_t is the **weighted** feature that is sent as input into one LSTM unit. For Example, If p_t is 0.6 then the image features (I_t) is weighted higher and the skeletal features (S_t) are weighted lower (1 - 0.6 = 0.4).

4. Experiments

4.1. Datasets

In skeleton-based action recognition the skeletal structure is represented as a graph. For our vision only system, we use state of the art 3D skeleton estimation LCR-Net [42] to estimate poses for the TUM Kitchen and UW-IOM dataset. While the focus of our work is to evaluate our proposed method on an online setting, we also run experiments on Skeleton Kinetics and NTU-RGB datasets as shown in Appendix Section I.

UW-IOM Dataset is a new dataset introduced in [35] with the intention of capturing activities that are common in warehouses. It consists of twenty videos (with average rate of twelve frames per second) of a sequence of object manipulation. The duration of every video is approximately three minutes. This dataset represents seventeen action classes and labels are of four-tier hierarchy indicating the object (box/rod), human motion (walk, stand, and bend), type of object manipulation if applicable (reach, pick-up, place, and hold), and the relative height of the surface where manipulation is taking place (low, medium, and high).

TUM Kitchen Dataset [53] consists of nineteen videos of a sequence of kitchen activities from four monocular cameras with the rate of twenty-five frames per second and the average duration of two minutes (we used camera 2). We use the provided two-tier labels by [35], which includes a motion verb (place, reach, stand), and a location (cabinet, drawer) or object manipulation mode (both-hands, one-hand) and creates a total of twenty-one activity classes.

4.2. Implementation Details

In our experiments, we sample a fixed length T=80 frames from each skeleton sequence as the input for online experiments. For offline experiments (NTU dataset and Skeleton Kinetics) we set the length T = 150 to cover the entire sequence for one label. We set the batch size to 128 and 32 for online and offline experiments respectively. In order to compare fairly with ST-GCN, the graph partitioning for the first adjacency matrix ($\hat{\mathbf{A}}_1$) is set to the same spatial strategy and partitioned into 3 subsets: the root node itself, centripetal group, and centrifugal group. However for the subsequent graphs $\hat{\mathbf{A}}_2$ and $\hat{\mathbf{A}}_3$ we assume that fully connected graph as initialization (all nodes are connected to every other node) and learn the edge importance weighting.

It should be noted that we do not modify the original ST-GCN model in terms of number of GCN or parameters. Our final model has only three GCN layers as opposed to the ten GCN-TCN components. More specifically the first GCN layer has 64 channels, second GCN has 128 and third has 256 channels. During training, we use the Adam optimizer [11] to optimize the network. We set the betas to 0.9 and 0.999 and set weight decay to zero. We split the training and validation using a **five** fold split in both TUM and UW-IOM. We report he mean and variance of all the splits in the results Table 2. We also do a grid search for learning rate(lr) from 0.1 to 0.001. On an average, lr of 0.05 performs best on all the splits in both datasets.

4.3. Ergonomics Risk Assessment

Given the input skeleton (\mathbf{X}) and the recognized activity from PGN (Figure 1), we compute REBA. REBA assigns human posture scores, in the range 1-15, based on joint angles during an activity. First, a risk score is computed for lower and upper extremities and those scores are added to the task-related scores (coupling and load scores). In [35] the score over all subjects are averaged offline and one score is reported for each activity class. We are proposing a realtime subject specific REBA evaluation using pose and action information such as the weight of the object and the type of manipulation.

5. Results and Discussion

5.1. Baseline Models

GCN vs Non-GCN Methods. To see the benefit of temporal analysis we perform experiments that only take human skeleton or image as input. We use these features as inputs of a *TCN* [1], *ED-TCN* [21] and *LSTM* [15] model. Baselines are trained in an online fashion. We also perform frame based experiments to determine the efficacy of temporal modelling. It must be noted that no additional convolution or linear layers are used to transform the input pose.

ST-GCN variants. We showcase the original ST-GCN implementations modified to support online setting by removing the final average pooling layer. Most ST-GCN variants used for spatio-temporal modelling, support recursive GCN-TCN models that pool messages across the overall graph of full skeleton. We also replace the 1x1 TCN convolutions with LSTMs. We refer to this model as *GCN+TCN*.

Modelities	Backhonos	UW-IOM			TUM		
wiouanties	Dackbolles	mAP (%)	Edit (%)	F1-overlap (%)	mAP (%)	Edit (%)	F1-overlap (%)
Skeleton (only)	Frame based	39.82 ± 1.45	29.26 ± 1.32	37.87 ± 1.82	29.79 ± 4.74	27.55 ± 2.89	32.63 ± 4.66
	LSTM [15]	79.35 ± 4.55	77.82 ± 6.34	85.32 ± 5.37	44.24 ± 5.97	56.46 ± 5.92	57.13 ± 8.24
	TCN [1]	57.72 ± 6.40	56.40 ± 5.36	64.78 ± 6.38	30.61 ± 5.40	51.07 ± 6.17	49.87 ± 11.01
	ED-TCN [21]	60.05 ± 4.89	81.73 ± 2.44	84.60 ± 2.64	28.89 ± 5.77	$\textbf{56.75} \pm \textbf{8.50}$	55.92 ± 11.11
	ST-GCN [60]	66.94 ± 3.49	61.89 ± 3.56	71.08 ± 2.83	34.73 ± 5.98	53.88 ± 5.53	53.52 ± 7.09
	ST-GCN+IMP [60]	73.28 ± 4.30	67.21 ± 6.05	76.58 ± 4.95	34.93 ± 4.75	52.27 ± 3.99	52.60 ± 5.72
	GCN+LSTM+IMP	81.97 ± 7.34	72.25 ± 7.24	82.04 ± 6.08	45.92 ± 4.19	52.07 ± 4.01	55.26 ± 5.54
	ST-PGN+LSTM (ours)	86.33 ± 2.71	77.92 ± 2.44	86.83 ± 1.74	48.02 ± 4.68	55.31 ± 5.09	57.58 ± 6.38
	ST-PGN+LSTM+IMP (ours)	85.92 ± 1.62	77.75 ± 2.46	86.21 ± 1.91	42.74 ± 1.03	47.19 ± 6.39	51.14 ± 6.94
	ST-PGN+LSTM+IMP+ML (ours)	$\textbf{87.03} \pm \textbf{2.85}$	$\textbf{97.86} \pm \textbf{2.15}$	$\textbf{87.95} \pm \textbf{1.54}$	$\textbf{49.62} \pm \textbf{6.10}$	56.10 ± 4.98	$\textbf{57.60} \pm \textbf{6.03}$
Image (only)	Frame based	51.62 ± 4.12	25.60 ± 1.55	34.17 ± 3.08	35.33 ± 5.26	28.33 ± 1.94	35.34 ± 2.65
	LSTM	66.50 ± 7.55	48.31 ± 5.90	57.81 ± 6.64	49.04 ± 7.03	52.64 ± 7.50	$\textbf{58.60} \pm \textbf{7.53}$
Fusion	Frame based+ Concat	50.54 ± 1.55	27.57 ± 0.96	36.42 ± 2.09	41.70 ± 5.76	29.66 ± 1.25	36.04 ± 1.59
	LSTM+ Concat	83.55 ± 5.74	72.98 ± 7.32	77.89 ± 11.70	48.71 ± 9.42	54.86 ± 6.83	57.11 ± 8.81
	ST-PGN+LSTM+IMP+ML+GRU-Fusion (ours)	$\textbf{87.05} \pm \textbf{3.47}$	80.90 ± 2.06	$\textbf{88.08} \pm \textbf{1.89}$	$\textbf{57.79} \pm \textbf{6.43}$	54.49 ± 5.59	58.35 ± 9.78

Table 2: mAP, edit, and F1-overlap score represented in mean and standard deviation over five splits in UW-IOM and TUM datasets for different methods and modalities. The best results in skeleton and fusion modality are shown in bold.

Edge Importance, as in the original work, is trained and is showcased as *ST-GCN+IMP*. LSTMs generally outperforms the TCNs to capture short transition changes in online fashion. Hence for the following experiments we choose to use LSTM as a primary temporal modelling source.

ST-PGN variants. Our models are showcased as pyramid-GCN (PGN) models. Similar to the previous section, we choose to train the edge importance for each of the sub graphs. The *predictions are averaged* for *ST-PGN+LSTM* and *ST-PGN+LSTM+IMP* and used to compute a single loss. Alternatively, our final multi-loss (ML) model has three losses, one for each of the pyramids. These losses are averaged and propagated during training. During testing the model's predictions are averaged and used for evaluation. The results for this model is shown in Table 2 under *ST-PGN+LSTM+IMP+ML*.

Fusion Models. To evaluate the impact of adding contextual features, we use a fusion mechanism that learns the importance of each feature modalities through a gauging mechanism (p_t in top-left of Figure 2).

5.2. Performance Analysis

The UW-IOM dataset focus is on object manipulation tasks that involve picking up, placing, and carrying objects, as well as walking bending and standing. Therefore, when we look at the edge importance demonstrations in the top left of Figure 3, we see that left hand (L-hand), right hand (R-hand) and right hip (R-hip) are the most important nodes in the low-level edge importance heat-map. Also, at the high-level, the importance of arms is higher than the legs and spine. We achieve an overall +5% improvement in mAP, +2% improvement in F1-overlap (*ST-PGN+LSTM*) over the best baseline (*GCN+LSTM* and *LSTM*). However, we see an overall performance boost of +16% in Edit score and similar to our multi-loss model. Importantly, ST-PGN is more powerful in distinguishing pick-up and place. These activities are spatially very similar and differ primarily in

temporal aspects. We do not see a huge benefit in Edit score using our image fusion. However, we see a minor improvement of 1% in the mAP and F1-overlap.

TUM kitchen dataset also includes object manipulation activities; however, it is focused on common daily activities in a kitchen. Looking at the low-level heat-map (top right in Figure 3) we observe that the hand, elbow, shoulder, and the neck joints have more importance. Looking at the high-level demonstration, we observe that the arms are more important than the legs and spine. We observe an overall improvement in mAP and F1-overlap using our models. However, a simple ED-TCN is slightly better at capturing the sequence and hence the Edit score is higher. Since the subjects move around in the scene, the significance of the lower body (*legs, hip*) is visibly higher in the edge importance compared to UW-IOM.

The results reported in Table 2 show that using skeleton is sufficient to get equal or better performance as compared to image-only or the fusion of skeleton and image. In UW-IOM, the human is facing the camera; thus, the detected skeleton is accurate. However, since this is not the case in the TUM dataset, the image-based models perform better as compared to the skeleton only on TUM dataset. If the skeleton is accurate, the addition of the image does not seem to enhance the results significantly in these tasks.

5.3. Failure Cases

We showcase confusion matrices of our best *ST*-*PGN+LSTM+IMP+ML* model in Figure 4 of the previous section. While we see an overall performance increase on both UW-IOM and TUM datasets, the model cannot deal with confusion among similar classes. We showcase the skeleton only model as adding image features do not help significantly. We describe our insights in detail below.

UW-IOM Dataset Our models can differentiate between *box-handling* actions and *rod-handling* actions without the use of image features (The skeleton configuration differs



Figure 3: Three level learned edge importance heat-map in UW-IOM (shaded) and TUM datasets. Each row shows the edge importance of each level of graph pyramid and it is consistent with bottom-left of Figure 2. Every level of PGN consists of the sum of three edge importance multiplied by the adjacency matrix and node features.

and handling of these objects is distinct due to the object size and location). However, *Standing* and *walking* misclassifications occur especially when the subject's back faces the camera. Hence important hand motions that help to infer these actions are missed. Better self-occlusion handling is warranted. Confusion also occurs between *bending* actions such as *bending-place*. This is predominantly due to misclassifications in transitions between these actions since bending is followed by pickup action or preceded by place action. Since it is a challenge for human annotators to accurately label transitions, the edit score should avoid penalizing such transitions.

TUM Dataset The camera view angle contributes to significant confusion between related classes. We choose the training-validation split with the lowest mAP score to analyze the results. The following observations are made:

1) The *pickup-drawer* and *close-drawer* are completely misclassified in this split. Once the drawer is closed, the pose estimation predicted the hand orientation and location using LCR-Net occlusion strategy [42]. However, the predicted pose is not always reliable, resulting in poor performance due to incorrect pose input during training. 2) *Walk-not holding* is misclassified for the majority of the



Figure 4: Confusion Matrix of *ST-PGN+LSTM+IMP+ML* model. Larger figures are added in the Appendix section.

classes such as *reach-cabinet*, *reach-drawer*, *stand-hold-both-hands*, *stand-not-hold*. This is attributed to unbalanced class distribution, where most of the actions are walking. In future work, we plan to address biases introduced by data imbalance by introducing sampling strategies. 3) *Twisting* actions are very challenging to detect using vision only since we only measure poses in Cartesian coordinates. Adding rotation information should help the model detect twisting actions about certain body axes. 4) *Pickup-hold-both-hands* gets confused with either *Pickup-hold-one-hand* or *stand-hold-both-hands*. Confusion is primarily due to one hand either being occluded by the object being handled, or the pose configuration being too similar in pose configuration with *standing*. More key-points in the pose prediction models could help resolve such issues.

6. Conclusion and Future Work

We proposed a novel Spatio-Temporal Pyramid Graph Convolutional Network (ST-PGN) for online action recognition. The method integrates the following: a) basic prior knowledge about the skeletal structure, b) hierarchical joint relationships and c) data-driven learning framework for online action based ergonomics risk assessment. The proposed approach addresses the simultaneous association of time-varying pose with action and objects interaction to enable downstream applications that involve computational modeling and prediction of various human performance metrics for ergonomics risk assessment.

Some open issues remain. First, generalization concerning other skeletal joint representations (Lie [56], Quaternion [36]) and camera viewpoint changes have not been addressed. Furthermore, while we outperform state of the art on online action recognition we are only comparable to the state of the art in offline action datasets (NTU-RGBD and Skeleton-Kinetics). The effect of the action label distribution on our models need to be studied further. In future work, we hope to address these issues with improved context fusion, long-term temporal modeling, and biomechanically consistent human pose representations [66].

References

- S. Bai, J. Z. Kolter, and V. Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.
- [2] P. P. Busto, A. Iqbal, and J. Gall. Open set domain adaptation for image and action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [3] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. arXiv preprint arXiv:1812.08008, 2018.
- [4] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [5] T. Cloete and C. Scheffer. Benchmarking of a full-body inertial motion capture system for clinical gait analysis. In 2008 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, pages 4579–4582. IEEE, 2008.
- [6] A. Colim, P. Carneiro, N. Costa, P. M. Arezes, and N. Sousa. Ergonomic assessment and workstation design in a furniture manufacturing industrya case study. In *Occupational and Environmental Safety and Health*, pages 409–417. Springer, 2019.
- [7] J. Dai, K. He, and J. Sun. Convolutional feature masking for joint object and stuff segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3992–4000, 2015.
- [8] R. Dey and F. M. Salemt. Gate-variants of gated recurrent unit (GRU) neural networks. In 2017 IEEE 60th international midwest symposium on circuits and systems (MWS-CAS), pages 1597–1600. IEEE, 2017.
- [9] Y. Du, W. Wang, and L. Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *Proceed*ings of the IEEE conference on computer vision and pattern recognition, pages 1110–1118, 2015.
- [10] X. Guo, J. Liu, and Y. Chen. Fitcoach: Virtual fitness coach empowered by wearable mobile devices. In *IEEE IN-FOCOM 2017-IEEE Conference on Computer Communications*, pages 1–9. IEEE, 2017.
- [11] I. Habibie, W. Xu, D. Mehta, G. Pons-Moll, and C. Theobalt. In the wild human pose estimation using explicit 2d features and intermediate 3d representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10905–10914, 2019.
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [13] S. Hignett and L. McAtamney. Rapid entire body assessment (reba). Applied ergonomics, 31(2):201–205, 2000.
- [14] S. Hignett and L. McAtamney. Rapid entire body assessment. In *Handbook of Human Factors and Ergonomics Methods*, pages 97–108. CRC Press, 2004.
- [15] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

- [16] H. Idrees, A. R. Zamir, Y.-G. Jiang, A. Gorban, I. Laptev, R. Sukthankar, and M. Shah. The THUMOS challenge on action recognition for videos in the wild. *Computer Vision* and Image Understanding, 155:1–23, 2017.
- [17] S. Kim, K. Yun, J. Park, and J. Y. Choi. Skeleton-based action recognition of people handling objects. In 2019 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 61–70. IEEE, 2019.
- [18] W. Kim et al. Adaptable workstations for human-robot collaboration: A reconfigurable framework for improving worker ergonomics and productivity. *IEEE Robot. Autom. Mag.*, pages 1–1, 2019.
- [19] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907, 2016.
- [20] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: a large video database for human motion recognition. In 2011 International Conference on Computer Vision, pages 2556–2563. IEEE, 2011.
- [21] C. Lea, M. D. Flynn, R. Vidal, A. Reiter, and G. D. Hager. Temporal convolutional networks for action segmentation and detection. In proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 156–165, 2017.
- [22] I. Lee, D. Kim, S. Kang, and S. Lee. Ensemble deep learning for skeleton-based action recognition using temporal sliding lstm networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1012–1020, 2017.
- [23] C. Li and S. Lee. Computer vision techniques for worker motion analysis to reduce musculoskeletal disorders in construction. In *Comput. Civil Eng.*, pages 380–387. 2011.
- [24] M. Li, S. Chen, X. Chen, Y. Zhang, Y. Wang, and Q. Tian. Actional-structural graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3595–3603, 2019.
- [25] Y. Li, C. Lan, J. Xing, W. Zeng, C. Yuan, and J. Liu. Online human action detection using joint classification-regression recurrent neural networks. In *European Conference on Computer Vision*, pages 203–220. Springer, 2016.
- [26] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision* and pattern recognition, pages 2117–2125, 2017.
- [27] C. Liu, Y. Hu, Y. Li, S. Song, and J. Liu. Pku-mmd: A large scale benchmark for continuous multi-modal human action understanding. arXiv preprint arXiv:1703.07475, 2017.
- [28] J. Liu, A. Shahroudy, G. Wang, L.-Y. Duan, and A. C. Kot. SSNet: scale selection network for online 3D action prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8349–8358, 2018.
- [29] J. Liu, G. Wang, P. Hu, L.-Y. Duan, and A. C. Kot. Global context-aware attention LSTM networks for 3D action recognition. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 1647– 1656, 2017.
- [30] A. Malaisé, P. Maurice, F. Colas, and S. Ivaldi. Activity recognition for ergonomics assessment of industrial tasks

with automatic feature selection. *IEEE Robotics and Automation Letters*, 4(2):1132–1139, 2019.

- [31] P. Maurice, A. Malaisé, C. Amiot, N. Paris, G.-J. Richard, O. Rochel, and S. Ivaldi. Human movement and ergonomics: An industry-oriented dataset for collaborative robotics. *The International Journal of Robotics Research*, 38(14):1529– 1537, 2019.
- [32] E. Mavroudi, D. Bhaskara, S. Sefati, H. Ali, and R. Vidal. End-to-end fine-grained action segmentation and recognition using conditional random field models and discriminative sparse coding. In 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 1558–1567. IEEE, 2018.
- [33] A. Mazzoldi, D. De Rossi, F. Lorussi, E. Scilingo, and R. Paradiso. Smart textiles for wearable motion capture systems. *AUTEX Research Journal*, 2(4):199–203, 2002.
- [34] S. Oh, A. Hoogs, A. Perera, N. Cuntoor, C.-C. Chen, J. T. Lee, S. Mukherjee, J. Aggarwal, H. Lee, L. Davis, et al. A large-scale benchmark dataset for event recognition in surveillance video. In *CVPR 2011*, pages 3153–3160. IEEE, 2011.
- [35] B. Parsa, E. U. Samani, R. Hendrix, C. Devine, S. M. Singh, S. Devasia, and A. G. Banerjee. Toward ergonomic risk prediction via segmentation of indoor object manipulation actions using spatiotemporal convolutional networks. *IEEE Robotics and Automation Letters*, 4(4):3153–3160, 2019.
- [36] D. Pavllo, D. Grangier, and M. Auli. Quaternet: A quaternion-based recurrent model for human motion. arXiv preprint arXiv:1805.06485, 2018.
- [37] L. Peppoloni, A. Filippeschi, E. Ruffaldi, and C. Avizzano. A novel wearable system for the online assessment of risk for biomechanical load in repetitive efforts. *International Journal of Industrial Ergonomics*, 52:1–11, 2016.
- [38] R. Poppe. A survey on vision-based human action recognition. *Image and vision computing*, 28(6):976–990, 2010.
- [39] G. Possebom, A. dos Santos Alonço, S. D. C. Bellochio, T. G. Lopes, D. P. Carpes, R. S. Becker, A. R. Moreira, T. R. Francetto, F. P. Rossato, and B. C. C. R. Zart. Comparison of methods for postural assessment in the operation of agricultural machinery. *Journal of Agricultural Science*, 10(9), 2018.
- [40] A. Prati, C. Shan, and K. I.-K. Wang. Sensors, vision and networks: From video surveillance to activity recognition and health monitoring. *Journal of Ambient Intelligence and Smart Environments*, 11(1):5–22, 2019.
- [41] G. Rogez and C. Schmid. Mocap-guided data augmentation for 3d pose estimation in the wild. In Advances in neural information processing systems, pages 3108–3116, 2016.
- [42] G. Rogez, P. Weinzaepfel, and C. Schmid. Lcr-net: Localization-classification-regression for human pose. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 3433–3441, 2017.
- [43] A. S. J. Roodbandi, F. Ekhlaspour, M. N. Takaloo, and S. Farokhipour. Prevalence of musculoskeletal disorders and posture assessment by qec and inter-rater agreement in this method in an automobile assembly factory: Iran-2016. In *Congress of the International Ergonomics Association*, pages 333–339. Springer, 2018.

- [44] K. Schaub, G. Caragnano, B. Britzke, and R. Bruder. The european assembly worksheet. *Theoretical Issues in Er*gonomics Science, 14(6):616–639, 2013.
- [45] A. Shafti, A. Ataka, B. U. Lazpita, A. Shiva, H. A. Wurdemann, and K. Althoefer. Real-time robot-assisted ergonomics. In *IEEE Int. Conf. Robot. Autom.*, 2019.
- [46] Z. Shou, D. Wang, and S.-F. Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1049–1058, 2016.
- [47] A. Shrivastava, A. Gupta, and R. Girshick. Training regionbased object detectors with online hard example mining. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 761–769, 2016.
- [48] C. Si, W. Chen, W. Wang, L. Wang, and T. Tan. An attention enhanced graph convolutional lstm network for skeletonbased action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1227–1236, 2019.
- [49] A. K. Singh, M. Meena, H. Chaudhary, and G. Dangayach. Ergonomic assessment and prevalence of musculoskeletal disorders among washer-men during carpet washing: guidelines to an effective sustainability in workstation design. *International Journal of Human Factors and Ergonomics*, 5(1):22–43, 2017.
- [50] G. Singh, S. Saha, M. Sapienza, P. H. Torr, and F. Cuzzolin. Online real-time multiple spatiotemporal action localisation and prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3637–3646, 2017.
- [51] 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.
- [52] W. Sultani and M. Shah. What If We Do Not Have Multiple Videos of the Same Action?–Video Action Localization Using Web Images. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 1077– 1085, 2016.
- [53] M. Tenorth, J. Bandouch, and M. Beetz. The TUM kitchen data set of everyday manipulation activities for motion tracking and action recognition. In *IEEE Int. Conf. Comput. Vis. Workshops*, pages 1089–1096, 2009.
- [54] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference* on computer vision, pages 4489–4497, 2015.
- [55] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference* on Computer Vision and Pattern Recognition, pages 6450– 6459, 2018.
- [56] R. Vemulapalli, F. Arrate, and R. Chellappa. Human action recognition by representing 3d skeletons as points in a lie group. In *Proceedings of the IEEE conference on computer* vision and pattern recognition, pages 588–595, 2014.
- [57] P. Wang, W. Li, P. Ogunbona, J. Wan, and S. Escalera. Rgb-d-based human motion recognition with deep learning: A survey. *Computer Vision and Image Understanding*, 171:118–139, 2018.

- [58] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A comprehensive survey on graph neural networks. arXiv preprint arXiv:1901.00596, 2019.
- [59] Z. Xu, Z. Liu, C. Sun, K. Murphy, W. T. Freeman, J. B. Tenenbaum, and J. Wu. Unsupervised Discovery of Parts, Structure, and Dynamics. 2018.
- [60] 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.
- [61] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2117–2126, 2017.
- [62] H. Zhao, M. Tian, S. Sun, J. Shao, J. Yan, S. Yi, X. Wang, and X. Tang. Spindle net: Person re-identification with human body region guided feature decomposition and fusion. In *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition, pages 1077–1085, 2017.
- [63] Y. Zhao, Y. Xiong, and D. Lin. Recognize actions by disentangling components of dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6566–6575, 2018.
- [64] W. Zheng, L. Li, Z. Zhang, Y. Huang, and L. Wang. Relational network for skeleton-based action recognition. In 2019 IEEE International Conference on Multimedia and Expo (ICME), pages 826–831. IEEE, 2019.
- [65] J. Zhou, G. Cui, Z. Zhang, C. Yang, Z. Liu, and M. Sun. Graph neural networks: A review of methods and applications. arXiv preprint arXiv:1812.08434, 2018.
- [66] Y. Zhu, B. Dariush, and K. Fujimura. Kinematic self retargeting: A framework for human pose estimation. *Computer* vision and image understanding, 114(12):1362–1375, 2010.