Multimodal multipart learning for action recognition in depth videos

Shahrudy, Amir; Ng, Tian-Tsong; Yang, Qingxiong; Wang, Gang

2016

Shahrudy, A., Ng, T.-T., Yang, Q., & Wang, G. (2016). Multimodal multipart learning for action recognition in depth videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(10), 2123-2129.

<https://hdl.handle.net/10356/87094>

<https://doi.org/10.1109/TPAMI.2015.2505295>

© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: [<http://dx.doi.org/10.1109/TPAMI.2015.2505295>].

Downloaded on 26 Apr 2023 17:35:12 SGT

Multimodal Multipart Learning for Action Recognition in Depth Videos

Amir Shahroudy, *Student Member, IEEE*,
Tian-Tsong Ng, *Member, IEEE*,
Qingxiong Yang, *Member, IEEE*,
and Gang Wang*, *Member, IEEE*

Abstract—The articulated and complex nature of human actions makes the task of action recognition difficult. One approach to handle this complexity is dividing it to the kinetics of body parts and analyzing the actions based on these partial descriptors. We propose a joint sparse regression based learning method which utilizes the structured sparsity to model each action as a combination of multimodal features from a sparse set of body parts. To represent dynamics and appearance of parts, we employ a heterogeneous set of depth and skeleton based features. The proper structure of multimodal multipart features are formulated into the learning framework via the proposed hierarchical mixed norm, to regularize the structured features of each part and to apply sparsity between them, in favor of a group feature selection. Our experimental results expose the effectiveness of the proposed learning method in which it outperforms other methods in all three tested datasets while saturating one of them by achieving perfect accuracy.

Index Terms—Action recognition, Kinect, Joint sparse regression, Mixed norms, Structured sparsity, Group feature selection.

1 INTRODUCTION

Human actions consist of simultaneous flow of different body parts. Based on this complex articulated essence of human movements, the analysis of these signals could be highly complicated. To ease the task of classification, actions could be broken down into their components. This is done by a body part detection on depth sequences of human body movements [1]. Having the 3D locations of body joints in the scene, we can separate the complicated motion of body into a concurrent set of behaviors on major skeleton joints; therefore human action sequences could be considered as multipart signals. Throughout this paper, we use the term “part” to denote each body joint as defined in [1].

Limiting the learning into skeleton based features cannot deliver high levels of performance in action recognition, because: (1) most of the usual human actions are defined based on the interaction of body with other objects, and (2) depth based skeleton data is not always accurate due to the noise and occlusion of body parts. To alleviate these issues, different depth based appearance features can be leveraged. The work in [2] proposed LOP (local occupancy patterns) around each of the body joints in order to represent 3D appearance of the interacting objects. Another solution is HON4D (histogram of

oriented 4D normals) [3], which gives more descriptive and robust models of the local depth based appearance and motion, around the joints. Based on the complementary properties of mentioned features, it is beneficial to utilize all of them as different descriptors for each joint. Combining heterogeneous features of each part of the skeleton, leads into a multimodal-multipart combination, which demands sophisticated fusion algorithms.

An interesting approach to handle the articulation of actions was recently proposed by [2]. As the key intuition, they have shown each individual action class can be represented by the behavior and appearance of few informative joints in the body. They utilized a data mining technique to find these discriminative sets of joints for each class of the available actions and tied up the features of those parts as “actionlets”. They employed a multi-kernel learning method to build up ensembles of actionlets as kernels for action classification. This method is highly robust against the noise in depth maps, and the results show its strength to characterize the human body motion and also human-object interactions. However the downside of this approach is the inconsistency of their heuristic selection process (mining actionlets) with the following learning step. Moreover, it simply concatenates different types of features for multimodal fusion, which is another drawback of this work. In this fashion, achieving the optimal combination of features regarding the classification task cannot be guaranteed.

To overcome the limitations mentioned above, we propose a joint structured sparsity regression based learning method which integrates part selection into the learning process considering the heterogeneity of features for each joint. We associate all the features for each part as a bundle and apply a group sparsity regularization to select a small number of active parts for each action class. To model the precise hierarchy of the multimodal-multipart features in an integrated learning and selection framework, we propose a hierarchical mixed norm which includes three levels of regularization over learning weights. To apply the modality based coupling over heterogeneous features of each part, it applies a mixed norm with two degrees of “diversity” induction [4], followed by a group sparsity among the feature groups of different parts to apply part selection.

The main contributions of this paper are two-fold: First, we integrated the part selection process into our learning in order to select discriminative body parts for different action classes latently, and utilize them to learn classifiers. Second, a hierarchical mixed norm is proposed to apply the desired simultaneous sparsity and regularization over different levels of learning weights corresponding to our special multimodal-multipart features in a joint group sparsity regression framework.

We evaluate our method on three challenging depth based action recognition datasets: MSR-DailyActivity dataset [2], MSR-Action3D dataset [5], and 3D-ActionPairs dataset [3]. Our experimental results show that the proposed method is superior to other available methods for action recognition on depth sequences.

The rest of this paper is organized as follows: Section 2 reviews the related works on depth based action recognition, joint sparse regression, mixed norms, and multitask learning. Section 3 presents the proposed integrated feature selection and learning scheme. It also introduces the new multimodal-multi part mixed norm which applies regularization and group sparsity into the proposed learning model. Experimental results on three above-mentioned benchmarks are covered in section 4 and we conclude the paper in section 5.

- A. Shahroudy and G. Wang are with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore, 639798.
E-mail: {amir3,wanggang}@ntu.edu.sg.
- A. Shahroudy and T.-T. Ng are with the Institute for Infocomm Research, 1 Fusionopolis Way, Singapore, 138632.
E-mail: {stusam,ttng}@i2r.a-star.edu.sg.
- Q. Yang is with the Department of Computer Science, City University of Hong Kong, Hong Kong.
E-mail: qiyang@cityu.edu.hk

* Corresponding author

Manuscript received 18 June 2014; revised 14 Mar. 2015; accepted 28 Oct. 2015.

2 RELATED WORK

Visual features extracted from depth signals can be classified into two major classes. The first are skeleton based features, which extract information from the provided 3D locations of body joints on each frame of the sequence. Essentially, skeletons have a very succinct and highly discriminative representation of the actions. [6] utilized them to extract “eigenjoints” for action classification using a Naïve-Bayes-Nearest-Neighbor classifier. In [7] spherical histograms of 3D locations of the joints went through HMM to model the temporal changes and final action classification. [8] split the body to groups of joint quads and represented each quad by a 6D vector. They utilized Fisher vectors to aggregate the quads as global representation of the action samples. Presence of noise in depth maps and occlusion of body parts bounds the reliability of this type of features. Another major deficiency of skeleton data is their incapacity to represent the interactions of the body with other objects which is crucial for activity interpretation.

The other group, consists of features which are extracted directly from depth maps. Most of the features in this class consider depth maps as spatio-temporal signals and tried to extract local or holistic descriptions from input sequences. [5] proposed a depth based action graph model in which each node indicates a salient posture and actions were represented as paths through graph nodes. To deal with occlusion and noise issues in depth maps, [9] proposed “random occupancy pattern” features and applied an elastic-net regularization [10] to find the most discriminative subset of features for action recognition. STIP (space-time interest point) detection described by HOG (histogram of oriented gradients) [11] and HOF (histogram of optical flow) was originally proposed for recognition purposes on RGB videos [12], but [13] showed this could be easily generalized into RGB+D signals. To improve the discrimination of descriptors, they generalized the idea of “motion history images” [14] over depth maps. Noise-suppression could also boost up the performance of STIP detection on depth sequences [15]. Four dimensional surface normals were shown to be very powerful representations of body movements over depth signals [3]. This idea was a generalization of HOG3D [16] into four dimensional depth videos. They quantized the 4D normal vectors of depth surfaces by taking their histograms over the vertices of a 4D regular polychoron, which were shown to be highly informative for action classification. In [17] local histograms of oriented gradients from spatio-temporal cells of action videos were aggregated by locality-constrained linear coding. Convolutional networks are utilized for depth based action recognition in [18]. Fusion of RGB+D modalities is also studied for action recognition [13], [19] and for other problems [20], [21].

Regarding the strengths and weaknesses of aforementioned classes of depth-based features, we infer they are complementary to each other and to achieve higher levels of performance, we have to combine them. [2] used histograms of 3D point clouds around the joints (LOP) to be added into skeleton based features for action classification using an “actionlet ensemble” framework. [22] added local HON4D [3] into joint features to learn a max-margin temporal warping based action classifier. We utilize skeletons, LOP and HON4D as state-of-the-art depth based features to build up our multimodal input for the task of action recognition.

The main intuition behind the work of [2] was the fact that features of few informative joints are good enough for recognizing each class of the actions. They defined “actionlet” as the combination of features of a limited numbers of joints

Fig. 1. Three Levels of the Proposed Hierarchical Mixed Norm for Multimodal Multipart Learning. We combine two levels of regularization inside modality groups and between them for each part, followed by a sparsity inducing norm between the parts to apply part selection.

and based on the discriminative power of each joint and each actionlet, they performed a data mining procedure to find the best actionlets for each class of the actions. They used mined actionlets as kernels in a multi-kernel multiclass SVM. We further extend this idea by applying group sparsity in a joint feature selection framework. To do so, we group the features of each part (joint) and applied L^1 norm between these groups to achieve a sparse set of active parts to represent each action class.

Mixed norms are powerful tools to inject simultaneous sparsity and coupling effects between the learning coefficients. They have been studied in a variety of fields. In statistical domain, [23] proposed the “group Lasso”, as an extension over “Lasso” [24] for a grouped variable selection in regression. [25] introduced “composite absolute penalty” for hierarchical variable selection. “Hierarchical penalization” is also proposed to utilize prior structure of the variables for a better fitting model [26]. In sparse regression, mixed norms have been used as regularization terms to link sparsity and persistence of variables [27]. A generalized shrinkage scheme was proposed by [28] for structured sparse regression. [29] used mixed norms as structured sparsity regularizers for heterogeneous feature fusion, and [30] extended this idea for a multi-view clustering. [31] proposed a robust self-taught learning using mixed norms and [32] utilized a fractional mixed norm for robust adaptive dictionary learning. In this paper, to regularize the multimodal features of each part, we apply a mixed L^2/L^4 norm. To achieve the sparsity between parts, we generalize this into an $L^1/L^2/L^4$ hierarchical norm.

If multiple learning tasks at hand share some inherent constituents or structures, “Multitask Learning” [33], [34], [35], [36] techniques could be globally beneficial. In joint sparse regression, multitask learning is formulated by a mixed norm. [37] proposed an L^1/L^∞ norm to add this into Lasso for variable selection. In joint feature selection, L^1/L^2 norm can provide multitask learning by applying selection between the L^2 regularized parameters of each feature [38]. Same is used in [39] as a generalization of L^1 norm in a multitask joint sparsity

representation model to fuse complementary visual features across recognition tasks. [40] studied different mixed norms when they applied multitask sparse learning in visual tracking and based on their experimental results, they showed L^1/L^2 is superior among them. In this work, we use a similar norm to utilize the shared latent factors between different binary action classifiers. We apply L^2 regularization over the weights corresponding to each feature across all the tasks, followed by an L^1 between all the features at hand.

3 MULTIMODAL MULTIPART LEARNING

Notations

Throughout this paper, we use bold uppercase letters to represent matrices and bold lowercase letters to indicate vectors. For a matrix $\mathbf{X}$, we denote its j -th row as $\mathbf{x}^j$ and its i -th column as $\mathbf{x}_i$.

Assume the partition ξ is defined over a vector $\mathbf{z}$ to divide its elements into $|\xi|$ disjoint sets. We use ξ_i to represent the indices of i -th set in ξ , and its corresponding elements in $\mathbf{z}$ are referred to as $\mathbf{z}^{\xi_i}$, also $z^{\xi_i, k}$ represents the k -th element of $\mathbf{z}^{\xi_i}$. The L^p/L^q norm of $\mathbf{z}$ regarding ξ is represented by $\|\mathbf{z}\|_{q,p|\xi}$ and is defined as the L^q norms of the elements inside each set of ξ followed by an L^p norm of the L^q values across the sets; mathematically:

$$\|\mathbf{z}\|_{q,p|\xi} = \left(\sum_{i=1}^{|\xi|} \|\mathbf{z}^{\xi_i}\|_q^p \right)^{1/p} = \left(\sum_{i=1}^{|\xi|} \left(\sum_{k=1}^{|\xi_i|} |z^{\xi_i, k}|^q \right)^{p/q} \right)^{1/p} \quad (1)$$

in which $|\xi_i|$ indicates the cardinality of set ξ_i .

Now consider the elements of each set ξ_i are further partitioned by operator ρ into $|\rho|$ disjoint subsets. Similarly, we indicate j -th ρ -subset of i -th ξ -set of $\mathbf{z}$ as $\mathbf{z}^{\xi_i, \rho_j}$ and $z^{\xi_i, \rho_j, k}$ represents its k -th element. The $L^p/L^q/L^r$ norm of $\mathbf{z}$ regarding ξ and ρ is also represented by $\|\mathbf{z}\|_{r,q,p|\rho,\xi}$ and is defined as the L^q/L^r norms (regarding ρ) of all $|\xi|$ sets followed by an L^p norm of the L^q/L^r values across the sets of ξ ; mathematically:

$$\begin{aligned} \|\mathbf{z}\|_{r,q,p|\rho,\xi} &= \left(\sum_{i=1}^{|\xi|} \|\mathbf{z}^{\xi_i}\|_{r,q}^p \right)^{1/p} \\ &= \left(\sum_{i=1}^{|\xi|} \left(\sum_{j=1}^{|\rho|} \left(\sum_{k=1}^{|\rho_j|} |z^{\xi_i, \rho_j, k}|^r \right)^{q/r} \right)^{p/q} \right)^{1/p} \quad (2) \end{aligned}$$

This representation can be easily extended into higher orders of structural mixed norms by further partitioning the subsets.

3.1 Multipart Learning by Structured Sparsity

Our purpose of learning is to recognize the actions in depth videos, based on depth based and skeleton based features extracted. The set of input features we use to describe each action sample is a combination of multimodal multipart features. The entire body is separated into a number of parts (as illustrated in Fig.1) and for each part we have different types of features to represent the movement and local depth appearance. Therefore, our input feature set for each input sample, can be represented by a vector: $\mathbf{z} \in \mathbb{R}^d$, which consists of feature groups of different parts and modalities. Assume operator π is partitioning $\mathbf{z}$ into P parts, and μ is defined over sets of π to further partition them based on M number of features modalities. So, the hierarchy of features inside this vector is indicated by: $\mathbf{z} = [\mathbf{z}^{\pi_1 T}, \dots, \mathbf{z}^{\pi_P T}]^T$, in which each $\mathbf{z}^{\pi_i} = [\mathbf{z}^{\mu_1, \pi_i T}, \dots, \mathbf{z}^{\mu_M, \pi_i T}]^T$.

Now the problem of multiclass action recognition can be considered as multiple binary regression based classification problems in a one versus all manner. Given n training samples $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ in which $\mathbf{x}_i \in \mathbb{R}^d$ and their corresponding labels for C distinct classes: $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_C]$ with $y_c \in \{0, 1\}^n$ and $\forall i : \sum_{c=1}^C y_c^i = 1$; we are looking for a projection matrix $\mathbf{W}^* \in \mathbb{R}^{d \times C}$ which minimizes a set of loss functions $J^c(\langle \mathbf{x}_i, \mathbf{w}_c^* \rangle, y_c^i)$ for all classes $c \in \{1, \dots, C\}$ and samples $i \in \{1, \dots, n\}$. Our choice for the total loss function, without loss of generality, is sum of squared errors ($\forall c : J^c(a, b) = (a - b)^2$).

The most common shrinkage methods to regularize the learning weights against overfitting are to penalize L^p norms of the learning weights for each class:

$$\mathbf{w}_c^* = \underset{\mathbf{w}_c}{\operatorname{argmin}} \sum_{i=1}^n J^c(\langle \mathbf{x}_i, \mathbf{w}_c \rangle, y_c^i) + \lambda \|\mathbf{w}_c\|_p \quad (3)$$

in which λ is the regularization factor. Employing L^2 norm ($p = 2$) leads into a general weight decay and minimization of the magnitude of $\mathbf{W}$, and applying L^1 norm ($p = 1$) yields simultaneous shrinkage and sparsity among the individual features. Such methods simply ignore the structural information between the features, which can be useful for classification; therefore, it is beneficial to embed these feature relations into our learning scheme via structured sparsity inducing mixed norms.

In the context of depth based action recognition, features are naturally partitioned into parts. "Actionlet ensemble" method [2] tried to discover discriminative joint groups using a data mining process, which led into an interesting improvement on the performance; however, their heuristic selection process is discrete and separated from the following learning step. To address these issues, we propose to apply group sparsity to perform part selection and classification in a regression based framework, in contrast to the mining based joint group discovery of [2].

We know that the discriminative strength of features in each part are highly correlated regarding all the classes at hand. So we expect the corresponding learning parameters (elements of each $\mathbf{w}_c$) to be triggered or halted concurrently within each set of π partitioning (for each action class). To apply a grouping effect on these features, we consider each set in π as a unit and measure its strength with an L^2 norm of the included learning weights. On the other hand, we seek a sparse set of parts to be activated for each class at hand, so we apply an L^1 norm between the L^2 values of the groups. Such an intuition can be formulated by an L^1/L^2 mixed norm based on π for each class:

$$\mathbf{w}_c^* = \underset{\mathbf{w}_c}{\operatorname{argmin}} \sum_{i=1}^n J^c(\langle \mathbf{x}_i, \mathbf{w}_c \rangle, y_c^i) + \lambda \|\mathbf{w}_c\|_{2,1|\pi} \quad (4)$$

Adding this up for all the action classes with the same regularization factor, we have:

$$\begin{aligned} \mathbf{W}^* &= \underset{\mathbf{W}}{\operatorname{argmin}} \sum_{c=1}^C \sum_{i=1}^n J^c(\langle \mathbf{x}_i, \mathbf{w}_c \rangle, y_c^i) + \lambda \sum_{c=1}^C \|\mathbf{w}_c\|_{2,1|\pi} \\ &= \underset{\mathbf{W}}{\operatorname{argmin}} \mathbf{J}(\mathbf{X}^T \mathbf{W}, \mathbf{Y}) + \lambda \|\operatorname{vec}(\mathbf{W})\|_{2,1,1|\pi,\tau} \quad (5) \end{aligned}$$

in which $\operatorname{vec}(\cdot)$ is the vectorization operator and τ is the partitioning operator of $\operatorname{vec}(\mathbf{W})$ elements based on their corresponding tasks (or columns here): $\forall (k, c) : \tau(w_c^k) = c$. We will refer to this multipart learning method as "MP".

Minimization of (5) applies the desired grouping effect into the features of each part and guarantees the sparsity on the number of active parts for each class in a smooth and simpler way, compared to the actionlet method.

3.2 Multimodal Multipart Learning via Hierarchical Mixed Norm

In the above formulation, we apply an L^2 regularization norm over heterogeneous features of all the modalities for each part, and ignore the modality structures between them. In other words, applying a general L^2 norm may cause the suppression of the information at some dimensions. These issues are more severe when training samples are limited (which is the case for action recognition in depth), in which it might lead to weak generalization of the learning.

To overcome these limitations, we utilize L^∞ to regularize the coefficients inside each modality, so that “diversity” [27] can be encouraged. It is already known that the behavior of L^p norm for $p > 2$ rapidly moves towards L^∞ [41]; since L^∞ is not easy to optimize directly, we picked L^4 as the most efficient approximation of it. Higher order norms like L^6 apply the same effect but with a slightly more expensive processing cost.

By applying the L^4 norm to regularize the weights in each modality group of each part, now we have a three-level $L^1/L^2/L^4$ mixed norm. Inner L^4 gives more “diversity” to regularize the features inside each partiality-modality subset. L^2 norm employs a magnitude based regularization over the L^4 values to link different modalities of each part, and the outer L^1 applies the soft part selection between the L^2/L^4 values of each action class (Fig.1).

Replacing the previous structured norm by the proposed hierarchical mixed norm in (5), we have:

$$\begin{aligned} \mathbf{W}^* &= \underset{\mathbf{W}}{\operatorname{argmin}} \mathbf{J}(\mathbf{X}^T \mathbf{W}, \mathbf{Y}) + \lambda \sum_{c=1}^C \|\mathbf{w}_c\|_{4,2,1|\mu,\pi} \\ &= \underset{\mathbf{W}}{\operatorname{argmin}} \mathbf{J}(\mathbf{X}^T \mathbf{W}, \mathbf{Y}) + \lambda \|\operatorname{vec}(\mathbf{W})\|_{4,2,1,1|\mu,\pi,\tau} \end{aligned} \quad (6)$$

here, π indicates the partitioning of features based on their source body part, and μ represents further partitioning of each part’s set regarding the modalities of the features. In the rest of this paper, we use the abbreviation “MMMP” to refer to this method. It is worthwhile to note, changing the inner norm to L^2 will reduce the hierarchical norm into a two level mixed norm, i.e. $\|\operatorname{vec}(\mathbf{W})\|_{2,2,1,1|\mu,\pi,\tau} = \|\operatorname{vec}(\mathbf{W})\|_{2,1,1|\pi,\tau}$ derived directly from the definition of hierarchical norm (2).

When different learning tasks have similar latent features, “Multitask Learning” [33] techniques can improve the performance of the entire system by applying information sharing between the tasks. Here we are learning classifiers for C different classes which essentially have lots of latent components in common, so pushing them to share some features is beneficial for the classification task. This can be done by applying an L^2 grouping on all the weights corresponding to each individual feature. Each of these L^2 values represents the magnitude of strength for its corresponding feature among all the tasks. Then applying an L^1 over the magnitudes can apply a shared variable selection considering all the tasks. Adding the new multitask term into (6), we have:

$$\begin{aligned} \mathbf{W}^* &= \underset{\mathbf{W}}{\operatorname{argmin}} \mathbf{J}(\mathbf{X}^T \mathbf{W}, \mathbf{Y}) + \lambda_1 \sum_{k=1}^d \|\mathbf{w}^k\|_2 \\ &\quad + \lambda_2 \|\operatorname{vec}(\mathbf{W})\|_{4,2,1,1|\mu,\pi,\tau} \end{aligned} \quad (7)$$

$$\begin{aligned} &= \underset{\mathbf{W}}{\operatorname{argmin}} \mathbf{J}(\mathbf{X}^T \mathbf{W}, \mathbf{Y}) + \lambda_1 \|\operatorname{vec}(\mathbf{W})\|_{2,1|\phi} \\ &\quad + \lambda_2 \|\operatorname{vec}(\mathbf{W})\|_{4,2,1,1|\mu,\pi,\tau} \end{aligned} \quad (8)$$

here, d is the number of rows in $\mathbf{W}$ which is equal to the size of the entire feature vector, and ϕ defines the partitioning of $\operatorname{vec}(\mathbf{W})$ elements based on their corresponding individual features: $\forall(k, c) : \phi(w_c^k) = k$.

Combining these two regularization terms can be considered as a trade off between sparsity and persistence of features [42] based on their relations across the parts, modalities, and between the action classes.

In our experiments, we use $P = 20$ body joints as partitioning operator π . Since each column of $\mathbf{W}$ has the same hierarchical partitioning as input features: $\mathbf{W} = [\mathbf{w}_c^j]$, in which c counts the number of classes and j counts the feature groups for P joints. The features for each joint come from $M = 3$ different modalities: skeletons, LOP, and HON4D; this defines the μ operator. Therefore, each $\mathbf{w}_c^j = [\mathbf{w}_c^{j,1T}, \dots, \mathbf{w}_c^{j,M^T}]^T$, in which each $\mathbf{w}_c^{j,m}$ is the corresponding weight elements to class c , joint j and modality m . This way (8) will be expanded to:

$$\begin{aligned} \mathbf{W}^* &= \underset{\mathbf{W}}{\operatorname{argmin}} \|\mathbf{X}^T \mathbf{W} - \mathbf{Y}\|_F^2 + \lambda_1 \sum_{k=1}^d \|\mathbf{w}^k\|_2 \\ &\quad + \lambda_2 \sum_{c=1}^C \sum_{j=1}^P \left(\sum_{m=1}^M \|\mathbf{w}_c^{j,m}\|_4^2 \right)^{1/2} \end{aligned} \quad (9)$$

3.3 Two Step Learning Approach

The downside of current formulation is the large number of weights to be learned simultaneously, compared to the size of training samples which are highly limited in current depth based action recognition benchmarks. To resolve this, we first learn the partially optimum weights for multipart features of each modality separately and then fine-tune them by the proposed multimodal multipart learning.

To learn the partially optimum weights for each modality m , we optimize:

$$\begin{aligned} \widehat{\mathbf{W}}_m &= \underset{\mathbf{W}_m}{\operatorname{argmin}} \mathbf{J}(\mathbf{X}_m^T \mathbf{W}_m, \mathbf{Y}) + \hat{\lambda}_1 \|\operatorname{vec}(\mathbf{W}_m)\|_{2,1|\phi} \\ &\quad + \hat{\lambda}_2 \|\operatorname{vec}(\mathbf{W}_m)\|_{2,1,1|\pi,\tau} \end{aligned} \quad (10)$$

After achieving the partially optimum point for each modality, we merge the $\widehat{\mathbf{W}}_m$ values for all M modalities:

$$\widehat{\mathbf{W}} = [\widehat{\mathbf{W}}_1^T, \dots, \widehat{\mathbf{W}}_M^T]^T \quad (11)$$

Next is to fine-tune the weights in the multimodal-multipart learning fashion, on a neighborhood of $\widehat{\mathbf{W}}$ values. To do so, we expect the global optimum weight not to diverge too much from their partially optimal points:

$$\begin{aligned} \mathbf{W}^* &= \underset{\mathbf{W}}{\operatorname{argmin}} \mathbf{J}(\mathbf{X}^T \mathbf{W}, \mathbf{Y}) + \lambda_1 \|\operatorname{vec}(\mathbf{W})\|_{2,1|\phi} \\ &\quad + \lambda_2 \|\operatorname{vec}(\mathbf{W})\|_{4,2,1,1|\mu,\pi,\tau} + \lambda_3 \|\mathbf{W} - \widehat{\mathbf{W}}\|_F^2 \end{aligned} \quad (12)$$

The last term in (12) will limit the deviation of learning weights from their partially optimal point, as we expect them to be just fine-tuned in this step.

Upon optimization over training data, the detection of the learned classifier for each testing sample $\mathbf{x}_i$ can be obtained by:

$$f(\mathbf{x}_i) = \underset{c}{\operatorname{argmax}} \langle \mathbf{x}_i, \mathbf{w}_c^* \rangle \quad (13)$$

The optimization steps are all done by “L-BFGS” algorithm using off-the-shelf “minFunc” tool [43].

4 EXPERIMENTS

This section describes our experimental setup details and then provides the results of the proposed method on three depth based action recognition benchmarks.

TABLE 1
Subject-wise Cross-Validation Performance Comparison of the Proposed Hierarchical Mixed Norm with Plain and Multipart Group Sparsity Norm on the MSR-DailyActivity Dataset

Method	Structure/Hierarchical Norm Used	Accuracy
L^2	$\ vec(\mathbf{W})\ _2^2$	80.61±2.49%
MP	$\ vec(\mathbf{W})\ _{2,1,1 \pi,\tau}$	81.55±2.43%
MMMP	$\ vec(\mathbf{W})\ _{4,2,1,1 \mu,\pi,\tau}$	84.03±2.16%

TABLE 2
Performance Comparison of the Proposed Method Using Plain/Structured/Hierarchical Norms on the Standard Evaluation Split of the MSR-DailyActivity Dataset

Method	Structure/Hierarchical Norm Used	Accuracy
L^1	$\ vec(\mathbf{W})\ _1$	86.88%
L^2	$\ vec(\mathbf{W})\ _2^2$	87.50%
MP	$\ vec(\mathbf{W})\ _{2,1,1 \pi,\tau}$	88.13%
MMMP	$\ vec(\mathbf{W})\ _{4,2,1,1 \mu,\pi,\tau}$	91.25%

4.1 Experimental Setup

All the provided experiments are done on Kinect based datasets. Kinect captures RGB frames, depth map signals and 3D locations of major joints. To have a fair comparison with other depth based methods, we ignore the RGB signals. Skeleton extraction is done automatically by Kinect's SDK based on the part-based human pose recognition system of [1]. On each frame, we have an estimation of 3D positions of 20 joints in the body. All of our features are defined based on these joints as the multipart partitioning operator (π); therefore, each feature necessarily belongs to one of these parts.

To represent skeleton based features, first we normalize the 3D locations of joints against size, position and direction of the body in the scene. This normalization step eases the task of comparison between body poses. On the other hand, the extracted body locations and directions could also be highly discriminative for some action classes like "walking" or "lying down"; therefore we add them into the features under a new auxiliary part. To encode the dynamics of skeleton based features, we apply "Fourier temporal pyramid" as suggested by [2] and keep first four frequency coefficients of each short time Fourier transformation. This leads into a feature vector of size 1,876 for each action sample.

In addition to skeleton based features, other modalities we use are local HON4D [3] and LOP [2] to represent depth based local dynamics and appearance around each joint. On each frame, LOPs are extracted on a (96,96,320)-sized depth neighborhood of each joint, which is divided into $3 \times 3 \times 4$ number of (32,32,80)-sized bins. To represent LOP based kinetics, we use a similar Fourier temporal pyramid transformation. HON4D features are also extracted locally over the location of joints on each frame. We encode HON4D features using LLC (locality-constrained linear coding) [44] to reduce their dimensionality while preserving the locality of 4D surface normals. Dictionary size of 100 is picked for the clustering step. LLC codes go through a max pooling over a 3 level temporal pyramid. Dimension of the features for LOP and HON4D are 5,040 and 14,000 respectively. The overall dimensionality of input features for each sample is 20,916.

TABLE 3
Performance Comparison on the Standard Evaluation Split of the MSR-DailyActivity Dataset using Single Modality and Multimodal Features.

Method	Modalities	Accuracy
Actionlet Ensemble [2]	LOP	61%
Proposed MP	LOP	79.38%
Orderlet Mining [45]	Skeleton	73.8%
Actionlet Ensemble [2]	Skeleton	74%
Proposed MP	Skeleton	79.38%
Local HON4D [3]	HON4D	80.00%
Proposed MP	HON4D	81.88%
Actionlet Ensemble [2]	Skeleton+LOP	85.75%
Proposed MMMP	Skeleton+LOP	88.13%
MMTW [22]	Skeleton+HON4D	88.75%
Proposed MMMP	Skeleton+HON4D	89.38%
DSTIP [15]	DCSF+LOP	88.20%
Proposed MMMP	Skeleton+LOP+HON4D	91.25%

4.2 MSR-DailyActivity3D Dataset

According to its intra-class variations and choices of action classes, MSR-DailyActivity dataset [2], is one of the most challenging benchmarks for action recognition in depth sequences. It contains RGB, depth, and skeleton information of 320 action samples, from 16 classes of daily activities in a living room. Each activity is done by 10 distinct subjects in two different ways and evaluations are applied over a fixed cross-subject setting; first five subjects are taken for training and others for testing. Unlike other datasets, MSR-DailyActivity has a more realistic variation within each class. Subjects used both hands randomly to do the activities, and samples of each class are captured in different poses.

First, to verify the strengths of our proposed hierarchical mixed norm, we evaluate the performance of the classification in a subject-wise cross-validation scenario. We evaluate the performance of the plain L^2 norm, the multipart structured norm (MP), and the proposed hierarchical mixed norm (MMMP), in all 252 possible train/test splits of 5 out of 10 subjects. To have a proper comparison between these norms, we have not applied the multitask term. The results of this experiment are shown in Table 1. Adding part based grouping, when it ignores the modality associations between the features, can slightly improve the performance from 80.61% into 81.55%. By adding multimodality grouping and applying the proposed hierarchical mixed norm, improvement is more significant and reaches 84.03%.

Next, we verify the results of our method by applying mentioned norms on the standard train/test split of the subjects. As provided in Table 2, applying simple feature selection using a plain L^1 norm leads into 86.88% of accuracy. By applying a plain L^2 norm on all the features we get 87.50%. Multipart learning regardless of heterogeneity of the modalities leads into 88.13%. Finally by adding the multipart learning via the proposed hierarchical mixed norm we reach the interesting accuracy of 91.25% on this dataset. Applying higher orders for the inner-most norm (like $L^1/L^2/L^6$) achieved the same level of accuracy at a slightly higher processing time.

To assess the strength of the proposed multipart learning, we evaluate our method on single modality setting using (10). As shown in Table 3, on skeleton based features, we got 79.38%

TABLE 4
Average Cross-Subject Performance for MSR-Action3D Dataset on Three Action Subsets of [5]

Method (protocol of [5])	Accuracy
Action Graph on Bag of 3D Points [5]	74.7%
Histogram of 3D Joints [7]	79.0%
EigenJoints [6]	83.3%
Random Occupancy Patterns [9]	86.5%
Skeletal Quads [8]	89.9%
HOG3D-LLC [17]	90.9%
Depth HOG [47]	91.6%
Lie Group [48]	92.5%
JAS+HOG ² [49]	94.8%
DL-GSGC+TPM [50]	96.7%
Proposed MMMP	98.2%

TABLE 5
Performance Comparison for MSR-Action3D Dataset Over All Action Classes

Method (protocol of [2])	Accuracy
Depth HOG [47] (as reported in [22])	85.5%
Actionlet Ensemble [2]	88.2%
HON4D [3]	88.9%
DSTIP [15]	89.3%
Lie Group [48]	89.5%
HOPC [51]	91.6%
Max Margin Time Warping [22]	92.7%
Proposed MMMP	93.1%

compared to 74% of the baseline actionlet method. Using LOPs, our method achieved 79.38% which is more than 18% higher than the actionlet's performance. For local HON4D features, we achieved 81.88% compared to 80.00% of the baseline local HON4D method. Now we use the partially learned weights of single modality multipart learning and employ them for the optimization of (12) to learn globally optimum projections. First we try the combination of skeleton based features with LOP. Using proposed learning, we get 88.13% of accuracy which outperforms the baseline's best result of 85.75%. [22] used skeleton and HON4D features in a temporal warping framework and got 88.75%. Our method outperforms it using the same set of features by achieving 89.38% of accuracy. And finally using all three modalities, our method achieves the performance level of 91.25%. Table 3 shows the complete set of results for this experiment.

Our implementation is done in MATLAB, and not fully optimized for time efficiency. The average training and testing time of MMMP on a 3.2 GHz Core-i5 machine are 170 and 2×10^{-4} seconds respectively, with no parallel processing.

It is worth pointing out some of the published works on this dataset applied other train/test splits, e.g. [46] reported 93.1% of accuracy on a leave-one-subject-out cross-validation. On this setup, proposed MMMP method achieves 97.5%.

4.3 MSR-Action3D Dataset

MSR-Action3D [5] is another depth based action dataset which provided depth sequences and skeleton information of 567 samples for 20 action classes. Actions are done by 10 different subjects, two or three times each. Evaluations are applied over

TABLE 6
Performance Comparison for 3D Action Pairs Dataset

Method	Accuracy
Depth HOG [47] (as reported in [22])	66.11%
Actionlet Ensemble [2] (as reported in [22])	82.22%
HON4D [3]	96.67%
Max Margin Time Warping [22]	97.22%
HOG3D-LLC [17]	98.33%
HOPC [51]	98.33%
Proposed MMMP	100.0%

another fixed cross-subject setting: Odd numbered subjects are taken for training and evens for testing. On one hand, depth sequences in this dataset have clean background which eases the recognition, and on the other hand, number of classes are higher than other datasets which could be a challenge for classification.

The reported results on this dataset are divided in two different scenarios. First is the average cross-subject performance on three action subsets defined in [5], and second is the overall cross-subject accuracy regardless of subsets, as done in [2]. Following [48], we call them as protocols of [5] and [2]. Tables 4 and 5 show the results. Although we still have the highest accuracy among the reported results, the achieved margin is not as large as other datasets. This is because of the simplicity of actions in this dataset. Since there is not any interaction with other objects, most of the classes are highly distinguishable using skeleton only features; therefore our multimodality could not boost up the results that much, but the multipart learning still shows its advantage over other methods.

4.4 3D Action Pairs Dataset

To emphasize the importance of the temporal order of body poses on the meaning of the actions, [3] proposed 3D Action Pairs dataset. It covers 6 pairs of similar actions. The only difference between each pair is their temporal order so they have similar skeleton, poses, and object shapes. Each action is performed by 10 subjects, 3 times. First five subjects are taken for testing and others for training. Based on the fewer number of the action classes and absence of intra-class variations, this is the easiest benchmark among depth based action recognition datasets and other methods already achieved very high accuracies on it.

Here we apply our full multimodal multipart learning method using all three available modalities of features. As shown in Table 6, the proposed method, outperforms all others and saturates the benchmark by achieving the perfect performance level on this dataset.

5 CONCLUSION

This paper presents a new multimodal multipart learning approach for action classification in depth sequences. We show that a sparse combination of multimodal part-based features can effectively and discriminatively represent all the available action classes at hand. Based on the nature of the problem, we utilize a heterogeneous set of features from skeleton based 3D joint trajectories, depth occupancy patterns and histograms of depth surface normals and show the proper way of using them as multimodal features set for each part.

The proposed method does the group feature selection, weight regularization, and classifier learning in a consistent

optimization step. It applies the proposed hierarchical mixed norm to model the proper structure of multimodal multi-part input features by applying a diversity norm over the coefficients of each part-modality group, linking different modalities of each part by a magnitude based norm, and utilizing a soft part selection by a sparsity inducing norm.

The provided experimental evaluations on three challenging depth based action recognition datasets show the proposed method can successfully apply the structure of the input features into a concurrent group feature selection and learning and confirm the strengths of the suggested framework compared to other methods.

ACKNOWLEDGMENTS

This research was carried out at the Rapid-Rich Object Search (ROSE) Lab at the Nanyang Technological University, Singapore. The ROSE Lab is supported by the National Research Foundation, Singapore, under its Interactive Digital Media (IDM) Strategic Research Programme.

The research is in part supported by Singapore Ministry of Education (MOE) Tier 2 ARC28/14, and Singapore A*STAR Science and Engineering Research Council PSF1321202099.

REFERENCES

- [1] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake, "Real-time human pose recognition in parts from single depth images," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [2] J. Wang, Z. Liu, Y. Wu, and J. Yuan, "Learning actionlet ensemble for 3d human action recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2014.
- [3] O. Oreifej and Z. Liu, "Hon4d: Histogram of oriented 4d normals for activity recognition from depth sequences," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [4] B. Rao and K. Kreutz-Delgado, "An affine scaling methodology for best basis selection," *IEEE Transactions on Signal Processing (TSP)*, 1999.
- [5] W. Li, Z. Zhang, and Z. Liu, "Action recognition based on a bag of 3d points," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops)*, 2010.
- [6] "Effective 3d action recognition using eigenjoints," *Journal of Visual Communication and Image Representation*, 2014.
- [7] L. Xia, C.-C. Chen, and J. Aggarwal, "View invariant human action recognition using histograms of 3d joints," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops)*, 2012.
- [8] G. Evangelidis, G. Singh, and R. Horaud, "Skeletal quads: Human action recognition using joint quadruples," in *International Conference on Pattern Recognition (ICPR)*, 2014.
- [9] J. Wang, Z. Liu, J. Chorowski, Z. Chen, and Y. Wu, "Robust 3d action recognition with random occupancy patterns," in *European Conference on Computer Vision (ECCV)*, 2012.
- [10] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2005.
- [11] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005.
- [12] I. Laptev and T. Lindeberg, "Space-time interest points," in *IEEE International Conference on Computer Vision (ICCV)*, 2003.
- [13] B. Ni, G. Wang, and P. Moulin, "Rgb-d-hudaact: A color-depth video database for human daily activity recognition," in *IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, 2011.
- [14] A. Bobick and J. Davis, "The recognition of human movement using temporal templates," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2001.
- [15] L. Xia and J. Aggarwal, "Spatio-temporal depth cuboid similarity feature for activity recognition using depth camera," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [16] A. Klaeser, M. Marszalek, and C. Schmid, "A spatio-temporal descriptor based on 3d-gradients," in *British Machine Vision Conference (BMVC)*, 2008.
- [17] H. Rahmani, A. Mahmood, D. Huynh, and A. Mian, "Action classification with locality-constrained linear coding," in *International Conference on Pattern Recognition (ICPR)*, 2014.
- [18] P. Wang, W. Li, Z. Gao, J. Zhang, C. Tang, and P. Ogunbona, "Action recognition from depth maps using deep convolutional neural networks," in *IEEE Transactions on Human Machine Systems (THMS)*, 2015.
- [19] A. Shahrourdy, G. Wang, and T.-T. Ng, "Multi-modal feature fusion for action recognition in rgb-d sequences," in *International Symposium on Communications, Control and Signal Processing (ISCCSP)*, 2014.
- [20] A. Wang, J. Lu, J. Cai, T.-J. Cham, and G. Wang, "Large-margin multi-modal deep learning for rgb-d object recognition," *IEEE Transactions on Multimedia (TMM)*, 2015.
- [21] A. Wang, J. Lu, J. Cai, G. Wang, and T.-J. Cham, "Unsupervised joint feature learning and encoding for rgb-d scene labeling," *IEEE Transactions on Image Processing (TIP)*, 2015.
- [22] J. Wang and Y. Wu, "Learning maximum margin temporal warping for action recognition," in *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [23] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2006.
- [24] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, 1996.
- [25] P. Zhao, G. Rocha, and B. Yu, "The composite absolute penalties family for grouped and hierarchical variable selection," *The Annals of Statistics*, 2009.
- [26] M. Szafrański, Y. Grandvalet, and P. Morizet-mahoudeaux, "Hierarchical penalization," in *Advances in Neural Information Processing Systems (NIPS)*, 2008.
- [27] M. Kowalski, "Sparse regression using mixed norms," *Applied and Computational Harmonic Analysis*, 2009.
- [28] M. Kowalski and B. Torrèsani, "Structured Sparsity: from Mixed Norms to Structured Shrinkage," in *Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, 2009.
- [29] H. Wang, F. Nie, H. Huang, and C. Ding, "Heterogeneous visual features fusion via sparse multimodal machine," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [30] H. Wang, F. Nie, and H. Huang, "Multi-view clustering and feature learning via structured sparsity," in *International Conference on Machine Learning (ICML)*, 2013.
- [31] H. Wang, F. Nie, and H. Huang, "Robust and discriminative self-taught learning," in *International Conference on Machine Learning (ICML)*, 2013.
- [32] H. Wang, F. Nie, W. Cai, and H. Huang, "Semi-supervised robust dictionary learning via efficient $l_{2,0+}$ -norms minimization," in *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [33] R. Caruana, "Multitask learning," *Machine Learning*, 1997.
- [34] Q. Zhou, G. Wang, K. Jia, and Q. Zhao, "Learning to share latent tasks for action recognition," in *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [35] L. Wang, N. T. Pham, T.-T. Ng, G. Wang, K. L. Chan, and K. Leman, "Learning deep features for multiple object tracking by using a multi-task learning strategy," in *IEEE International Conference on Image Processing (ICIP)*, 2014.
- [36] A. Abdulnabi, G. Wang, J. Lu, and K. Jia, "Multi-task cnn model for attribute prediction," *IEEE Transactions on Multimedia (TMM)*, 2015.
- [37] H. Liu, M. Palatucci, and J. Zhang, "Blockwise coordinate descent procedures for the multi-task lasso, with applications to neural semantic basis discovery," in *International Conference on Machine Learning (ICML)*, 2009.
- [38] G. Obozinski, B. Taskar, and M. I. Jordan, "Joint covariate selection and joint subspace selection for multiple classification problems," *Statistics and Computing*, 2010.
- [39] X.-T. Yuan, X. Liu, and S. Yan, "Visual classification with multitask joint sparse representation," *IEEE Transactions on Image Processing (TIP)*, 2012.
- [40] T. Zhang, B. Ghanem, S. Liu, and N. Ahuja, "Robust visual tracking via multi-task sparse learning," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.

- [41] A. Rakotomamonjy, R. Flamary, G. Gasso, and S. Canu, " $l_p - l_q$ penalty for sparse linear and sparse multiple kernel multitask learning," *IEEE Transactions on Neural Networks*, 2011.
- [42] M. Kowalski and B. Torrsani, "Sparsity and persistence: mixed norms provide simple signal models with dependent coefficients," *Signal, Image and Video Processing*, 2009.
- [43] M. Schmidt, "Minfunc," 2005. [Online]. Available: <http://www.di.ens.fr/~mschmidt/Software/minFunc.html>
- [44] J. Wang, J. Yang, K. Yu, F. Lv, T. Huang, and Y. Gong, "Locality-constrained linear coding for image classification," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [45] G. Yu, Z. Liu, and J. Yuan, "Discriminative orderlet mining for real-time recognition of human-object interaction," in *Asian Conference on Computer Vision (ACCV)*, 2014.
- [46] S. Althloothi, M. H. Mahoor, X. Zhang, and R. M. Voyles, "Human activity recognition using multi-features and multiple kernel learning," *Pattern Recognition*, 2014.
- [47] X. Yang, C. Zhang, and Y. Tian, "Recognizing actions using depth motion maps-based histograms of oriented gradients," in *ACM International Conference on Multimedia (MM)*, 2012.
- [48] R. Vemulapalli, F. Arrate, and R. Chellappa, "Human action recognition by representing 3d skeletons as points in a lie group," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [49] E. Ohn-Bar and M. Trivedi, "Joint angles similarities and hog² for action recognition," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops)*, 2013.
- [50] J. Luo, W. Wang, and H. Qi, "Group sparsity and geometry constrained dictionary learning for action recognition from depth maps," in *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [51] H. Rahmani, A. Mahmood, D. Q. Huynh, and A. Mian, "Hopc: Histogram of oriented principal components of 3d pointclouds for action recognition," in *European Conference on Computer Vision (ECCV)*, 2014.