

Unsupervised Video Understanding by Reconciliation of Posture Similarities

Timo Milbich, Miguel Bautista, Ekaterina Sutter, Björn Ommer Heidelberg Collaboratory for Image Processing IWR, Heidelberg University, Germany

{timo.milbich, miguel.bautista, ekaterina.sutter, björn.ommer}@iwr.uni-heidelberg.de

Abstract

Understanding human activity and being able to explain it in detail surpasses mere action classification by far in both complexity and value. The challenge is thus to describe an activity on the basis of its most fundamental constituents, the individual postures and their distinctive transitions. Supervised learning of such a fine-grained representation based on elementary poses is very tedious and does not scale. Therefore, we propose a completely unsupervised deep learning procedure based solely on video sequences, which starts from scratch without requiring pre-trained networks, predefined body models, or keypoints. A combinatorial sequence matching algorithm proposes relations between frames from subsets of the training data, while a CNN is reconciling the transitivity conflicts of the different subsets to learn a single concerted pose embedding despite changes in appearance across sequences. Without any manual annotation, the model learns a structured representation of postures and their temporal development. The model not only enables retrieval of similar postures but also temporal super-resolution. Additionally, based on a recurrent formulation, next frames can be synthesized.

1. Introduction

The ability to understand human actions is of cardinal importance for our interaction with another. Explaining the activity of another person by observing only individual postures and their temporal transitions in a sequence of video frames has been a long-standing challenge in Computer Vision. There are numerous applications in problems like activity indexing and search [42, 38, 27], action prediction [22, 45], behavior understanding and transfer [35, 33, 7], abnormality detection [3], and action synthesis and video generation [12, 46, 4, 37]. Our goal is to learn human activity on the finest accessible level, i.e., individual poses, by capturing characteristic postures and the distinctive transitions between them solely based on videos without requiring any manual supervision or pre-defined body models. The underlying deep learning approach is unsupervised



Figure 1. Visualizing all frames of all long-jump sequences using the learnt posture representation ϕ . Similar frames across different sequences and repetitions across time are mapped nearby, yielding a concise rendering of the overall activity with its characteristic gait cycles. Moreover, intermediate frames can be synthesized (top right) as well as future frames of a sequence (visualized by nearest neighbors from training set (right)).

and starts from scratch with only the video data and without requiring tedious user input or pre-trained networks. Our approach to activity understanding, summarized in Fig. 1 (cf. Supplementary¹),exhibits the following characteristics:

Unsupervised: The most prominent paradigm to video understanding has been supervised action classification [22], since labeling finer entities such as individual poses [1] is tedious. Action classification typically utilizes wholistic models and a discriminative approach is trained to classify actions into discrete classes. As a result, such approaches model actions globally in terms of their overall most salient differences, rather than capturing the subtle changes of human posture over time (the clothing of a person may sufficie to discriminate running from diving). Model-free: Multiple works have modeled activity and pos-

¹Video material demonstrating the different applications can be found under https://hciweb.iwr.uni-heidelberg.de/compvis/ research/tmilbich_iccv17

ture using a predefined model for joint locations (i.e. Mo-Cap [40, 21, 54], Depth [36], etc.). However, obtaining this meta information is costly and prevents scaling these approaches to use large unlabeled collections of video data. Continuous in time: Several approaches have tackled the problem of understanding activity by decomposing it into discrete sub-actions [26, 53, 47] or into a hierarchy [54, 43]. Consequently, the detailed, continuous evolution between consecutive postures is neglected. Multi-granular: A lot of work has focussed separately either on pose matching [29, 5, 44], action classification [22, 48, 13, 2], or mid-level entities (e.g. clusters of postures) [53, 43, 26]. Explaining activities, however, demands to describe overall activity based on fine-granular postures and the transitions in between, thus linking coarse with fine granularities. Finegrained activity parsing: Explaining activity on the temporal scale of single postures with all their diverse changes is far more detailed and complex than mere action classification [48, 13]. Previous efforts [17] have approximated posture and its transitions using discrete states in an AND-OR graph and relied on tedious supervision information.

With no supervision information, no predefined model for human posture, and training from scratch, we need to compensate for there being no labels for individual postures. Since we are lacking the labels to directly train a representation for posture we can only utilize a large number of video frames and reason about pairwise relationships between postures. This is aggravated by the fact that the visual representation of posture in different videos can be significantly different due to changes in lighting, background, or the clothing and skin color of different persons. Therefore, we utilize a large number of training video frames and have a deep learning algorithm alternate between proposing pairwise similarities/dissimilarities between postures, and then resolving transitivity conflicts to bring the relationships into mutual agreement. To propose similarities, a combinatorial sequence matching algorithm is presented, which can find exact solutions, but only for small sets of frames. A CNN then resolves the transitivity conflicts between these different subsets of the training data by learning a posture embedding that reconciles the pairwise constraints from the different subsets. While the sequence matching is already incorporating information about posture changes, a Recurrent Neural Network (RNN) is trained to capture the overall activity and to predict future frames of an activity sequence by synthesizing transitions.

Experimental results show that our approach is able to successfully explain an activity by understanding how posture continuously changes over time and to model the temporal relationships between postures. Furthermore, our posture representation obtains state-of-the-art results on the problem of zero-shot human pose estimation and has also proven worth as a powerful initialization for other supervised human pose estimation methods. In addition, our approach captures the temporal progression of an activity, it can predict future frames, and it also enables a Generative Adversarial Network (GAN) to provide temporal superresolution.

2. Representation Learning for Parsing Activities

2.1. Learning a Posture Embedding

We are interested in detailed understanding of human activity without requiring manual interaction or predefined models. Therefore, we can explain overall activity in a video only using the most basic entity that we can directly access: the human posture observed in bounding box detections I of individual frames. Activity, which emerges at the temporal scale of an entire video sequence, is then represented by individual poses and their characteristic transitions and repetitions on a fine temporal scale. To model an activity, we need a posture representation $\phi(I; \theta)$ which: (i) is invariant to changes in environmental conditions such as lighting and background. (ii) is invariant to the appearance of persons (clothing and skin color). (iii) is continuous in time (consecutive frames are near in feature space). Only then we can understand the essential characteristics of an activity and spot all repetitions of the same pose over time and in different sequences, despite changes in person appearance or environment.

A natural choice to incarnate $\phi(I;\theta)$ are CNNs. They exhibit great expressive power to learn highly non-linear representations at the cost of requiring millions of manually labeled samples for training. A popular alternative is then to only fine-tune a representation that has been pre-trained for discrete classification on large datasets [8]. Unfortunately, the performance of pre-trained models for transfer learning is heavily task dependent. In our scenario, the discrete classification objective contradicts our requirements for ϕ , i.e., a discrete classification loss neglects smoothness within the representation space ϕ . In addition, since datasets like [8] are composed of single images, pre-trained models fail to encode temporal relationships between frames. This explains the inferior performance of these pre-trained models in Sect. 3.

Instead of using a pre-trained representation we seek to learn a representation $\phi(I; \theta)$ that maps similar postures close in feature space while retaining temporal structure of an activity. Ideally, manual supervision of human postures, such as joint annotations, and/or positive links of similar postures within and across sequences together with negative links of dissimilar postures could be used to learn $\phi(I; \theta)$ employing, for example, triplets of similar and dissimilar poses. However, we are lacking these labels, altogether. To overcome this lack of labels we exploit the relationships in-



Figure 2. Correspondences for two query sequences obtained using nearest neighbors (first row) and our sequence matching with temporal constraints (second row). Note how the temporal constraints provide much more accurate correspondences.

herent in large collections of video sequences. We infer the supervision information, which is required to learn a CNN representation $\phi(I; \theta)$, by solving a combinatorial sequence matching problem. The solution of this matching problem then provides us with correspondences of similar and dissimilar postures, which we then impose onto the CNN representation $\phi(I; \theta)$ to learn it.

2.2. Sequence Matching for Self-supervision

We aim to learn a CNN representation $\phi(I; \theta)$ which encodes posture similarity, without being provided with any labels. In order to learn such a representation, we employ a self-supervision strategy, leveraging the temporal information in videos to solve a sequence matching problem and find pair-wise correspondences between frames on a sequence level. Let $S = \{I_j\}_{j=1}^n$ and $S' = \{I'_{j'}\}_{j'=1}^{n'}$ denote two sequences of n and n' frames respectively, we want to find a correspondence π : $\{1, ..., n\} \mapsto \{0, 1, ..., n'\}$ that matches frames of S to frames of S', where the index 0 is used to match outliers. Furthermore, in order to successfully learn $\phi(I;\theta)$ using self-supervision, we want to enforce the following constraints on π : (i) Corresponding frames should be similar in appearance. (ii) To avoid temporal cross-over and to reduce false positives matches, consecutive correspondences must be chronologically ordered. (iii) To avoid only part of a sequence being used and to explore the full span of possible postures therein, one-to-many correspondences should be penalized. (iv) Correspondences should be invariant to the sequence frame rate. Figure 2

demonstrates the need for our temporal constraints. These constraints prevent, by definition, to utilize classical sequence matching approaches like the computational costly String Matching [24] or Dynamic Time Warping [6, 23]. Thus, we define the following optimization problem which combines all these constraints,

$$\begin{array}{l} \underset{\pi:\{1,\dots,n\}\mapsto\{0,1,\dots,n'\}}{\min} \sum_{j=1}^{n} \left\| \phi(I_{j};\theta) - \phi(I_{\pi(j)}';\theta) \right\|_{2}^{2} \\ + \lambda_{1} \sum_{j=1}^{n-1} \mathbb{1}_{\pi(j)>\pi(j+1)} + \lambda_{2} \sum_{j=1}^{n-1} \mathbb{1}_{\pi(j)=\pi(j+1)} \quad , \quad (1) \\ + \lambda_{3} \sum_{j=1}^{n-1} \mathbb{1}_{\pi(j)+1<\pi(j+1)} \left[\pi(j+1) - \pi(j) \right] \end{array}$$

where $\mathbb{1}(\cdot)$ denotes the indicator function and $\lambda_1, \lambda_2, \lambda_3$ penalize the violations of the different temporal constraints. The use of inequalities ensures constraint *(iv)*. To solve the optimization problem in Eq. (1) we convert it into an Integer Linear Program (ILP). In order to do so, we define a matrix $\mathbf{Z} \in \{0, 1\}^{n \times n' \times n'}$, where $z_{j,j'_1,j'_2} := \mathbb{1}_{\pi(j)=j'_1 \wedge \pi(j+1)=j'_2}$. A non-zero *z* indicates matches for two consecutive frames starting at position *j*. The ILP is then

$$\max_{\mathbf{Z} \in \{0,1\}^{n \times n' \times n'}} \sum_{j=1}^{n-1} \sum_{j_1', j_2'=0}^{n'} z_{j,j_1',j_2'} p_{j,j_1',j_2'}$$

subject to
$$\sum_{j_1',j_2'=0}^{n'} z_{j,j_1',j_2'} = 1$$

$$\wedge \sum_{j_1'=0}^{n'} z_{j,j_1',j_2'} = \sum_{j_3'=0}^{n'} z_{j+1,j_2',j_3'}$$
 (2)

where p_{j,j'_1,j'_2} is the sum of all terms in Eq. (1) with $z_{j,j'_1,j'_2} = 1$. To obtain reliable self-supervision information, we obtain an exact solution of this ILP using a branch-and-cut algorithm [34]. However, exactly solving this problem for pairs of long sequences (e.g. n > 500) is a costly operation with exponential worst case complexity in n, making it computationally infeasible.

To circumvent this high cost when matching S onto S', we break the target sequence S' into k equal length subsequences of length $n' \approx 40$ and find an exact solution for Eq. (2) on each local sub-sequence in parallel. Thus, the overall computational cost is reduced by a factor of k, making it a feasible to tackle long sequences. However we obtain only local correspondences at sub-sequence level and thus, discarding important relationships between different sub-sequences that compose the overall activity. To compensate for this shortcoming, we train a CNN with the different, local sub-sequence solutions in subsequent minibatches as discussed in Sect. 2.3. The CNN then reconciles the local sub-sequence correspondences. Thus, we benefit from combining the computational feasibility of exact local sub-sequence matching with the power of stochastic CNN training, which aggregates lots of local observations in one concerted representation.

2.3. From Local Correspondences to a Globally Consistent Posture Representation

Our training procedure combines multiple exact solutions to local ILPs to obtain self-supervision for training a joint CNN representation $\phi(I;\theta)$ that reconciles all the sub-problems. Since our goal is to learn a CNN representation for encoding human activity in a fully unsupervised manner, the mini-batches for training are composed just by pairs of sequences $\{S, S'\}$. We find these pairs by first randomly choosing S and sampling S' from a set of nearest neighbour sequences S_{nn} to S, thus sorting out totally unrelated sequences. S_{nn} is constructed using simple sequence descriptors by temporally pooling similarities over all frames of a video. After breaking S' into equallength sub-sequences, the ILP in Eq. (2) yields a solution \mathbf{Z}^* . These are exact pair-wise correspondences between \mathcal{S} and a particular sub-sequence of S'. We then use these correspondences \mathbf{Z}^* to generate triplets \mathcal{T}_t using a triplet sampling $\{\mathcal{S}, \mathcal{S}', \mathbf{Z}^*\} \mapsto \{\mathcal{T}_t, \}_{t=1}^T$. Here $\mathcal{T}_t = \{I_a, I_+, I_-\}$ consists of a randomly sampled anchor image $I_a \in \mathcal{S}$ and its positive correspondence $I_+ = \pi(I_a) \in S'$, together with a randomly chosen negative $I_{-} \in S'$. We randomly sample negatives based on the p-th percentile of the similarity distribution of sequence S' to I_+ . That is, we compute the similarity of I_+ to each frame of S', and sample negatives from frames with a lower similarity than the *p*-th percentile. We thus include hard negatives by decreasing pover epochs. Note that by sampling positives and negatives from the same sequence and by comparing the same sequence with different other sequences, relationships within sequences are also implicitly established. Using this triplet self-supervision we update the CNN parameters θ via backpropagation and the triplet ranking loss [50, 49].

$$L(\{\mathcal{T}_{t}\}_{t=1}^{T};\theta) = \frac{1}{T} \sum_{t=1}^{T} L'(\mathcal{T}_{t};\theta)$$
(3)
$$L'(\mathcal{T}_{t};\theta) = \left[\|\phi(I_{a};\theta) - \phi(I_{+};\theta)\|_{2}^{2} - \|\phi(I_{a};\theta) - \phi(I_{-};\theta)\|_{2}^{2} + \delta \right]_{+}$$
(4)

with δ controlling the margin between I_+ and I_- with respect to I_a . The matching algorithm from Eq. (2) provides the self-supervision information needed to train the CNN representation ϕ , whereas the CNN training using Eq. (3) yields the posture embedding required to compute the similarities in Eq. (2). Alg. 1 outlines this iterative procedure. We found that using HOG-LDA [18] to initialize



Figure 3. Similarity matrices for a long jump sequence of Olympic Sports dataset computed using (a) VGG-S pre-trained on Imagenet, (b) CliqueCNN [5], (c) Ours. The diagonal structures correspond to repetitions of gait cycles during running.

 ϕ for the first epoch provides a decent initialization leading to a speed-up compared to a random ϕ . Learning a single posture representation that captures characteristic similarities across and within sequences is thus decomposed into a series of mini-batch optimizations. For each, Eq. (2) provides a matching that is locally, within the respective sub-sequences, optimal. The stochastic optimization of the CNN then consolidates, over a number of mini-batches, the transitivity conflicts between the local solutions to arrive at a single posture representation. That way, both approaches combine their strengths and weaknesses in an ideal manner. Fig.3(a-c) show an excerpt of similarity matrices from different models. Note the significantly improved signal-tonoise ratio in (c).

2.4. RNN for Learning Temporal Transitions

We now have an algorithm which effectively learns an overall posture representation $\phi(I; \theta)$ for an activity based on relationships within and across sequences. This detailed posture encoding is an ideal basis to also deal with coarser temporal scales, enabling to learn the complete temporal structure of an activity without requiring any prior model and to synthesize future frames of a sequence. We employ an RNN to explain or synthesize overall activity based on the transitions between individual postures. As we are living in a continuous state space with our representation ϕ , also the transitions will be continuous. RNNs naturally lend themselves to encode temporally encoded entities, but significant modeling effort can be required to obtain

Algorithm 1	: Unsupervised	l learning o	of a consistent		
posture representation using local correspondences.					

```
 \begin{array}{l} \textbf{Data: } \{\mathcal{S}_i\}_{i=1}^{S}, \theta_{s=0} \\ // \text{ Unlabeled video sequences and randomly} \\ \text{ initialized } \theta_{s=0} \\ \textbf{Result: } \{\phi, \theta\} \\ \textbf{while } \|\theta_{s+1} - \theta_s\|_2 > \epsilon \ \textbf{do} \\ & \left(\mathcal{S}, \mathcal{S}'\right) \leftarrow \{\mathcal{S}_i\}_{i=1}^S // \text{ Training batch} \\ \textbf{Z}^* \leftarrow & \arg\min \quad \text{Eq. (2) } (\mathcal{S}, \mathcal{S}', \theta) \\ & \textbf{Z} \in \{0, 1\}^{n \times n' \times n'} \\ & \left\{\mathcal{T}_t\}_{t=1}^T \leftarrow \{\mathcal{S}, \mathcal{S}', \mathbf{Z}^*\} // \text{ Sample triplets} \\ & \theta_{s+1} \leftarrow \theta_s + \alpha \nabla_{\theta_s} L(\{\mathcal{T}_t\}_{t=1}^T; \theta_s) // \text{ Update } \theta \end{array} \right.
```



Figure 4. Top 5 LSTM predictions (red) for the next frame t + 1, given four previous frames (green). Actual successor in blue.

a meaningful representation. Thus intensive pre-training of the static entities (such as words or objects) is used before learning transitions between them (sentences, videos) [28].

However, the already structured posture space that we have learned in Sect. 2.3 simplifies this and we can incarnate our RNN as a standard LSTM [16], which takes as input $\phi(I; \theta)$. In order for the LSTM to be able to model an activity as the temporal transition between individual postures, we formulate a regression task for future pose prediction under representation $\phi(I; \theta)$. By training the LSTM to predict the representation of future frames based on the observation of a small preceding sub-sequence, the network learns to model the evolution of the complete activity. Let $\mathcal{C}^{t,l} = \{I_{t-l}, \ldots, I_t\} \subset \mathcal{S}$ be an l frame sub-sequence of S. The LSTM hidden state vector $\mathbf{h}_t \in \mathbb{R}^m$ has $m \gg d$ so as to effectively learn the variability of temporal transitions (cf. Sect. 3). We then stack one fully connected layer on top of h to revert to the original dimensionality d, thus generating a representation $\phi'(\phi(I_t, \theta); \theta') \in \mathbb{R}^d$ (a stacking of an LSTM layer and a fully-connected layer). By formulating next frame prediction as a regression problem, we train this network using the euclidean loss function

$$L_{\text{RNN}}(\mathcal{C}^{t,l}, I_{t+1}, \phi, \theta; \theta') = \|\phi'(\phi(I_t, \theta); \theta') - \phi(I_{t+1}; \theta)\|_2^2$$
(5)

The trained LSTM is then able to hypothesize transitions to future frames based on a small subsequence of an activity as demonstrated in Fig. 4.

3. Experimental Evaluation

To evaluate our approach at all granularity levels of an activity, we report quantitative results for posture retrieval in Sect. 3.1 and human pose estimation (HPE) in Sects. 3.2-3.3. In addition, we provide qualitative results for activity understanding in Sect. 3.4, and for temporal super-

0.62	0.56	0.64	0.58
Shuffle&Learn [31]	CliqueCNN [5]	Ours Scratch	Ours Imagenet
0.63	0.79	0.83	0.83

Table 1. Avg. AUC for each method on Olympic Sports dataset.

resolution, and action synthesis in Sects. 3.5 to 3.7.

3.1. Posture Retrieval: Olympic Sports dataset

The Olympic Sports (OS) dataset [32] is a compilation of video sequences of 16 different sports, containing more than 170000 frames overall in 300 video sequences. During training, each training mini-batch is generated by sampling two random sequences S, S' and solving the ILP in Eq. (2) to obtain correspondences \mathbf{Z}^* . Solving the ILP takes \sim 0.1sec (IBM CPLEX Optimization framework) compared to ~ 0.5sec to process a minibatch T_t on a NVIDIA Titan X (Pascal). We then use these correspondences to sample $\{\mathcal{T}_t,\}_{t=0}^{T=300}$ triplets, where the initial percentile p = 100 is decreased by 10 every epoch. In each frame the approach of [14] yields person bounding boxes. During training we utilize VGG-S up to the fc6 layer and stack a 128-dimensional fc7 layer on top together with an l_2 -normalization layer, which is our representation $\phi(I;\theta)$. We use *Caffe* [19] for our implementation.

To evaluate our representation on fine-grained posture retrieval we utilize the annotations provided by [5] and follow their evaluation protocol, using their annotations only for testing. We compare our method with CliqueCNN [5], the triplet formulation of Shuffle&Learn [31], the tuple approach of Doersch et. al [9], VGG-S [41], and HOG-LDA [18]. For completeness we also include a version of our model that was initialized with Imagenet pre-trained weights [41]. (i) For CliqueCNN, Shuffle& Learn, and Doersch et. al methods we use the models downloaded from their respective project websites. (ii) Exemplar-CNN is trained using the best performing parameters reported in [11] and the 64c5-128c5-256c5-512f architecture. Then we use the output of fc4 and compute 4-quadrant max pooling. During training of our approach, each image in the training set is augmented by performing random translation, scaling and rotation to improve invariance.

In Tab. 1 we show the average AuC over all categories for the different methods. When compared with the best method so far [5], the proposed approach improves the performance by 4%, although the method in [5] was even pretrained on Imagenet. This improvement is due to the cross sequence relationships enforced by our sequence matching, which enforce a representation which is invariant to background and environmental factors, encoding only posture. In addition, when compared to the state-of-the-art methods that leverage tuples [9] or triplets [31] for training a CNN from scratch, our approach shows 20% higher performance. This is explained by the more detailed similarity relationships encoded in the cross-sequence correspondences obtained by the sequence matching approach, which uses temporal constraints to obtain high quality relationships of similarity and dissimilarity. It is noteworthy, that our randomly initialized VGG-S trained with our self-supervision strategy yields equivalent performance to a version with pre-trained Imagenet weights as initialization. Thus, the proposed selfsupervision circumvents the use of the 1.2M labelled Imagenet samples. To the best of our knowledge, this is the first time that a self-supervised method performs equivalently without the widely adopted Imagenet pre-training strategy.

3.2. Zero-Shot HPE

After evaluating the proposed method for fine-grained posture retrieval, we tackle the problem of zero-shot pose estimation on the LSP dataset. That is, we transfer the pose representation learnt on Olympic Sports to the LSP dataset without any further training and retrieve similar poses based on their similarity. The LSP [20] dataset is one of the most widely used benchmarks for pose estimation. For evaluation we use the representation to compute visual similarities and find nearest neighbours to a query frame. Since the evaluation is zero-shot, joint labels are not available. At test time we therefore estimate the joint coordinates of a query person by finding the most similar frame from the training set and taking its joint coordinates. We then compare our method with VGG-S [41] pre-trained on Imagenet, the triplet approach of Misra et. al (Shuffle&Learn) [31] and CliqueCNN [5]. In addition, we also report an upper bound on the performance that can be achieved by zero-shot evaluation using ground-truth similarities. Here the most similar pose for a query is given by the frame which is closest in average distance of ground-truth pose annotations. This is the best one can achieve without a parametric model for pose (the performance gap to 100% shows the discrepancy between poses in test and train set). For completeness, we compare with a fully supervised state-of-the-art approach for pose estimation [51]. For computing similarities we now use the intermediate pool5 layer of VGG-S as our representation $\phi(I;\theta)$, provided that our model is transferred from another dataset [52].

In Tab. 2 we show the PCP@0.5 obtained by the different methods. For a fair comparison with CliqueCNN [5] (which was pre-trained on Imagenet), we include a version of our method trained using Imagenet initialization. Since in this experiment we are transferring our model from another dataset, we expect that Imagenet pre-training increases performance. Our approach significantly improves the visual similarities learned using both Imagenet pretrained VGG-S and CliqueCNN [5], obtaining a performance boost of at least 3% in PCP score. In addition, when trained from scratch without any pre-training on Imagenet our model outperforms the model of [31] by 13%, due to the fact that the cross-sequence correspondences obtained by

813						
01.5	54.6	48.8	36.1	19.1	56.9	50.0
80.1	50.1	45.7	27.2	12.6	45.5	43.5
82.0	48.2	41.8	32.4	15.8	53.6	47.0
73.0	45.1	41.6	26.2	12.2	44.4	43.0
60.4	33.2	28.9	16.8	7.1	33.8	30.0
93.7	78.8	74.9	58.7	36.4	72.4	69.2
98.4	95.0	92.8	88.5	81.2	95.7	90.9
	80.1 82.0 73.0 60.4 93.7 98.4	80.1 50.1 82.0 48.2 73.0 45.1 60.4 33.2 93.7 78.8 98.4 95.0	80.1 50.1 45.7 82.0 48.2 41.8 73.0 45.1 41.6 60.4 33.2 28.9 93.7 78.8 74.9 98.4 95.0 92.8 6 1 41.4	80.1 50.1 45.7 27.2 82.0 48.2 41.8 32.4 73.0 45.1 41.6 26.2 60.4 33.2 28.9 16.8 93.7 78.8 74.9 58.7 98.4 95.0 92.8 88.5	80.1 50.1 45.7 27.2 12.6 82.0 48.2 41.8 32.4 15.8 73.0 45.1 41.6 26.2 12.2 60.4 33.2 28.9 16.8 7.1 93.7 78.8 74.9 58.7 36.4 98.4 95.0 92.8 88.5 81.2	80.1 50.1 45.7 27.2 12.6 45.5 82.0 48.2 41.8 32.4 15.8 53.6 73.0 45.1 41.6 26.2 12.2 44.4 60.4 33.2 28.9 16.8 7.1 33.8 93.7 78.8 74.9 58.7 36.4 72.4 98.4 95.0 92.8 88.5 81.2 95.7

Table 2. PCP measure for each method on Leeds Sports dataset for zero-shot pose estimation.

our sequence matching approach encode finer relationships between samples. Finally, it is notable that even though our pose representation is *transferred from a different dataset* without fine-tuning on LSP, it obtains state-of-the-art performance in the realm of unsupervised methods.

3.3. Self-supervision as Pre-training for HPE

In addition to the zero-shot learning experiment we also evaluate our approach on the challenging MPII Pose dataset [1] which is a state of the art benchmark for evaluation of articulated human pose estimation. The dataset includes around 25K images containing over 40K people with annotated body joints. MPII Pose is a particularly challenging dataset because of the clutter, occlusion and number of persons appearing in images. We are interested in evaluating how far our self-supervised approach can still boost a parametric approach that is trained with extensive supervision. Thus, we report the performance obtained by DeepPose [44], when trained using as initialization each of the following models: Random initialization, Shuffle&Learn [31], Imagenet and our approach trained on OS (scratch and Imagenet pretraining). For this experiment the Alexnet [25] architecture is used like in [44]. Following the standard evaluation metric on MPII dataset, Tab. 3 shows the PCKh@0.5 obtained by training DeepPose (stg-1) using their best reported parameters with the different initializations.

The performance obtained on MPII Pose benchmark shows that our unsupervised feature representation successfully scales to challenging datasets, successfully dealing with clutter, occlusions and multiple persons. In particular, when comparing our unsupervised initialization with a random initialization we obtain a 5.1% performance boost, which indicates that our features encode a robust notion of pose that is robust to the clutter present in MPII dataset. Furthermore, we obtain a 1.2% improvement over the Shuffle&Learn [31] approach.

3.4. Visualizing the Activity Representation

Whereas previous experiments have evaluated our representation on the level of individual poses, we now analyze the ability of $\phi(I; \theta)$ to also capture transitions between successive postures. Our model is trained using temporally aligned correspondences across sequences, thus representing not only relationships between different sequences



Figure 5. Visualizing the learned posture representation ϕ and the progression of an activity (indicated by color). *All* frames of *all* vault sequences are shown (a) Our representation successfully learns the inherent structure of an action, e.g. repetitive gait and spinning cycles (blue, orange loops). Also, on a coarser temporal scale, repeated postures are brought near (outstretched arms before/after jump; cyan and red). Vector quantization of ϕ yields mutually dissimilar, characteristic poses shown in frames. (b) Representation obtained using [31]. As this model discards cross-sequence interactions, it misses the regularity of related postures across time and sequences.

but also encoding temporal transitions of pose within a sequence. Therefore, ϕ captures all the regularity of posture. It maps similar postures of different persons and repetitions of the same posture in a sequence, e.g., repeated gait cycles, to the same spot in feature space. Moreover, successive poses are also mapped to similar representations, so an activity has a smooth trajectory over ϕ . To visually demonstrate the ability of $\phi(I;\theta)$ to capture the finegrained pose interactions over time and between sequences

	Ours	S&I	Rand	Imagenet	Ours + Im-
	scratch	[31]	Init.	imagenet	agenet
Head	80.6	75.8	79.5	87.2	90.2
Neck	88.4	86.3	87.1	93.2	93.8
LR Shoulder	74.8	75.0	71.6	85.2	86.3
LR Elbow	56.9	59.2	52.1	69.6	70.4
LR Wrist	41.6	42.2	34.6	52.0	58.6
LR Hip	73.3	73.3	64.1	81.3	82.4
LR Knee	63.6	63.1	58.3	69.7	73.2
LR Ankle	56.9	51.7	51.2	62.0	67.4
Thorax	88.6	87.1	85.5	93.4	93.7
Pelvis	79.9	79.5	70.1	86.6	88.4
Total	70.5	60.3	65.4	78.0	80.4

Table 3. PCKh@0.5 measure on MPII Pose benchmark dataset using different initializations for the DeepPose approach [44].

we project the high dimensional representation ϕ to a 2D plot using the t-SNE procedure [30]. Fig. 5(a) shows a mapping of all instances of vault activity. Successive postures within each video are connected by straight lines and color encodes the time within the sequence. The learned representation captures the repetitive structure of running and spinning (blue and orange loops) and the characteristic transitions between. Additionally, the regularity of ϕ allows to group repetitive postures and to provide a condensed overview of an activity. Therefore, we employ standard agglomerative clustering [15] to extract prominent mutually dissimilar posture that span an activity. We show the representative of each cluster on its corresponding location in Fig. 5(a). Moreover, we compare our representation with the state-of-the-art approach of Misra et al. [31] which introduces a temporal verification problem and learns to find the correct temporal order of triplets of postures within a sequence, Fig. 5(b). This shows that modeling only posture interactions within a sequence [31] and excluding crosssequences correspondences as proposed in Sect. 2 fails to capture the temporal evolution of an activity and degrades temporal structure.

3.5. Inferring Temporal Super-Resolution

The previous section has demonstrated that our representation $\phi(I; \theta)$ successfully encodes posture and provides a basis for modeling the characteristic progression of an activity on the finest accessible level—individual frames. To further demonstrate the fine granularity at which activity is captured, we now unfold transitions between postures in consecutive frames, that is, we obtain super-resolution between two consecutive frames. Whereas [39] aggregates frames from within the same sequence, we bring all the different sequences with all their variability together. Since our representation maps related poses close to another, we can employ a local linear interpolation between successive postures to infer intermediate frames. Then the Generative Adversarial Network (GAN) of [10] is used to invert the feature back to an image. We use the implementation



Figure 6. Inferring intermediate frames between consecutive postures. For each transition we interpolate representations ϕ at regular intervals in between and invert interpolated features with [10].

of [10] for both the generator and discriminator networks, where our learned activity representation $\phi(I;\theta)$ acts as the encoder network. To jointly train the three networks we use the DeePSiM-loss [10] considering adversarial and euclidean terms on both the image and feature domain. This inversion of our representation $\phi(I;\theta)$ creates images for the synthesized intermediate frames, allowing us to go past the limited temporal scale of given video sequences. Fig. 6 shows temporal super-resolution results for two different activities. The continuity of our pose representation $\phi(I;\theta)$. It has finer temporal granularity than an individual video, since it interleaves a large number of related sequences, providing a truly continuous activity representation.

3.6. Activity Understanding using LSTMs

So far we have provided a comprehensive analysis of our activity representation, demonstrating its ability to understand actions on the fine-grained scale of single postures and beyond. Now, we evaluate the capability of our method to understand activity at the sequence level. Therefore, we train a recurrent network on top of the posture representation ϕ to yield a sequence-level encoding ϕ' , as discussed in Sect. 2.4. We employ an LSTM, trained on sub-sequences $C^{t,l}$ of length l = 4, sampled densely from all video sequences to predict the next succeeding frame I_{t+1} . During training, we sample mini-batches to cover the overall diversity of activity, so that all constituent postures are equally represented for learning the LSTM. Fig. 4 shows exemplary predictions from different activities.

Let us quantify the quality of predicted next frames. Given the true successor frame I_{t+1} , we identify its nearest neighbor in all the videos. We then compute the distance between these two frames and average it over all videos. The same is then done for the second, third, etc. nearest neighbor. Similarly, we compute the distance of our prediction and the true I_{t+1} and also average it. Fig. 7 (b) compares



Figure 7. (a) Setup of quantitative evaluation. Blue is the actual next frame, red the prediction, purple the nearest neighbours. (b) LSTM evaluation: Comparing our average prediction error (red) against the average distance of I_{t+1} to each of its 10 nearest neighbours (blue). The error bars indicate the standard deviation of the measurements.



Figure 8. Synthesizing an activity by recursively predicting the next posture. Green: initial image . Every 5th frame is shown.

the resulting error of our prediction against that of the k nearearest neighbor from the dataset. Our prediction is better than actually observing the next frame and picking its second nearest neighbor. Despite the large variability of an activity this shows that the temporal progression of an activity has been well captured to yield favorable predictions of a successive frame.

3.7. Video Understanding by Action synthesis

We have just seen predictions of the next frame I_{t+1} of a sequence. By recursively adding this predicted frame and then predicting a next successive frame we can iteratively synthesize an overall activity frame by frame. For visualization of the predicted next posture, we choose the nearest neighbor from the training set. Fig. 8 summarizes the synthesis of a snatch activity initialized at the green posture. One can see that our model successfully infers the temporal ordering of the activity from its beginning until the end.

4. Conclusions

In this paper we have presented an unsupervised approach for understanding activity by means of its most fine-grained temporal constituents, individual human postures. A combinatorial sequence matching algorithm proposes relations between frames from subsets of the training set, which a CNN uses to learn a single concerted pose embedding that reconciles transitivity conflicts. Without any manual annotation, the model learns a structured representation of postures and their temporal development. The model not only enables retrieval of similar postures but also temporal super-resolution. Additionally, based on a recurrent formulation, future frames and activities can be synthesized.

Acknowledgements: This work has been supported in part by the Heidelberg Academy for the Sciences. We are grateful to the NVIDIA corporation for donating a Titan X GPU.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recogni tion (CVPR)*, June 2014. 1, 6
- [2] B. Antic and Björn Ommer. Learning latent constituents for recognition of group activities in video. In Proceedings of the European Conference on Computer Vision, 2014. 2
- [3] Borislav Antić and Björn Ommer. Video parsing for abnormality detection. In *Computer Vision (ICCV)*, 2011 IEEE International Conference on, pages 2415– 2422. IEEE, 2011. 1
- [4] Borislav. Antic and Björn Ommer. Per-sample kernel adaptation for visual recognition and grouping. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015. 1
- [5] Miguel A Bautista, Artsiom Sanakoyeu, Ekaterina Sutter, and Björn Ommer. Cliquecnn: Deep unsupervised exemplar learning. *NIPS*, 2016. 2, 4, 5, 6
- [6] Donald J Berndt and James Clifford. Using dynamic time warping to find patterns in time series. In *KDD workshop*, volume 10, pages 359–370. Seattle, WA, 1994. 3
- [7] B. Brattoli, U. Büchler, A. S. Wahl, M. E. Schwab, and Björn Ommer. Lstm self-supervision for detailed behavior analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2017. 1
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE, 2009. 2
- [9] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, pages 1422–1430, 2015. 5
- [10] A. Dosovitskiy and T. Brox. Generating images with perceptual similarity metrics based on deep networks. In Advances in Neural Information Processing Systems (NIPS), 2016. 7, 8
- [11] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *NIPS*, pages 766–774, 2014. 5

- [12] Alexei A Efros, Alexander C Berg, Greg Mori, Jitendra Malik, et al. Recognizing action at a distance. In *ICCV*, volume 3, pages 726–733, 2003. 1
- [13] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *The IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), June 2016. 2
- [14] Pedro Felzenszwalb, David McAllester, and Deva Ramanan. A discriminatively trained, multiscale, deformable part model. In *CVPR*, pages 1–8. IEEE, 2008. 5
- [15] Hichem Frigui and Raghu Krishnapuram. Clustering by competitive agglomeration. *Pattern recognition*, 30(7):1109–1119, 1997. 7
- [16] Felix A Gers, Jürgen Schmidhuber, and Fred Cummins. Learning to forget: Continual prediction with lstm. *Neural computation*, 12(10):2451–2471, 2000.
 5
- [17] Abhinav Gupta, Praveen Srinivasan, Jianbo Shi, and Larry S Davis. Understanding videos, constructing plots learning a visually grounded storyline model from annotated videos. In *Computer vision and pattern recognition*, 2009. CVPR 2009. IEEE conference on, pages 2012–2019. IEEE, 2009. 2
- [18] Bharath Hariharan, Jitendra Malik, and Deva Ramanan. Discriminative decorrelation for clustering and classification. In *ECCV*, pages 459–472. Springer, 2012. 4, 5
- [19] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. arXiv preprint arXiv:1408.5093, 2014. 5
- [20] Sam Johnson and Mark Everingham. Learning effective human pose estimation from inaccurate annotation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2011. 6
- [21] Asha Kapur, Ajay Kapur, Naznin Virji-Babul, George Tzanetakis, and Peter F Driessen. Gesture-based affective computing on motion capture data. In *International Conference on Affective Computing and Intelligent Interaction*, pages 1–7. Springer, 2005. 2
- [22] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional

neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. 1, 2

- [23] Eamonn Keogh. Exact indexing of dynamic time warping. In *Proceedings of the 28th international conference on Very Large Data Bases*, pages 406–417.
 VLDB Endowment, 2002. 3
- [24] Donald E Knuth, James H Morris, Jr, and Vaughan R Pratt. Fast pattern matching in strings. *SIAM journal* on computing, 6(2):323–350, 1977. 3
- [25] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012. 6
- [26] Tian Lan, Yuke Zhu, Amir Roshan Zamir, and Silvio Savarese. Action recognition by hierarchical mid-level action elements. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4552– 4560, 2015. 2
- [27] Ivan Laptev, Marcin Marszalek, Cordelia Schmid, and Benjamin Rozenfeld. Learning realistic human actions from movies. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008. 1
- [28] Thang Luong, Richard Socher, and Christopher D Manning. Better word representations with recursive neural networks for morphology. In *CoNLL*, pages 104–113, 2013. 5
- [29] Fengjun Lv and Ramakant Nevatia. Single view human action recognition using key pose matching and viterbi path searching. In 2007 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2007), 18-23 June 2007, Minneapolis, Minnesota, USA, 2007. 2
- [30] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008. 7
- [31] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, pages 527–544. Springer, 2016. 5, 6, 7
- [32] Juan Carlos Niebles, Chih-Wei Chen, and Li Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. In *ECCV*, pages 392–405. Springer, 2010. 5

- [33] Nuria M Oliver, Barbara Rosario, and Alex P Pentland. A bayesian computer vision system for modeling human interactions. *IEEE transactions on pattern analysis and machine intelligence*, 22(8):831– 843, 2000. 1
- [34] Manfred Padberg and Giovanni Rinaldi. A branchand-cut algorithm for the resolution of large-scale symmetric traveling salesman problems. *SIAM review*, 33(1):60–100, 1991. 3
- [35] Maja Pantic, Alex Pentland, Anton Nijholt, and Thomas S Huang. Human computing and machine understanding of human behavior: A survey. In *Artifical Intelligence for Human Computing*, pages 47–71. Springer, 2007. 1
- [36] Miguel Reyes, Gabriel Dom'inguez, and Sergio Escalera. Featureweighting in dynamic timewarping for gesture recognition in depth data. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 1182–1188. IEEE, 2011.
- [37] J. C. Rubio, A. Eigenstetter, and Björn Ommer. Generative regularization with latent topics for discriminative object recognition. *Pattern Recognition*, 2015.
 1
- [38] Sreemanananth Sadanand and Jason J Corso. Action bank: A high-level representation of activity in video. In *Computer Vision and Pattern Recognition (CVPR)*, 2012 IEEE Conference on, pages 1234–1241. IEEE, 2012. 1
- [39] Oded Shahar, Alon Faktor, and Michal Irani. Superresolution from a single video. In *CVPR*, 2011. 7
- [40] Leonid Sigal, Alexandru O Balan, and Michael J Black. Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International journal of computer vision*, 87(1-2):4, 2010. 2
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 5, 6
- [42] Stephen W Smoliar and HongJiang Zhang. Content based video indexing and retrieval. *IEEE multimedia*, 1(2):62–72, 1994. 1
- [43] Bing Su, Jiahuan Zhou, Xiaoqing Ding, Hao Wang, and Ying Wu. Hierarchical dynamic parsing and encoding for action recognition. In *European Conference on Computer Vision*, pages 202–217. Springer, 2016. 2

- [44] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1653–1660, 2014. 2, 6, 7
- [45] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating visual representations from unlabeled video. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 98–106, 2016. 1
- [46] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In Advances In Neural Information Processing Systems, pages 613–621, 2016. 1
- [47] Jacob Walker, Abhinav Gupta, and Martial Hebert. Patch to the future: Unsupervised visual prediction. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 3302–3309, 2014. 2
- [48] Heng Wang, A. Klaser, C. Schmid, and Cheng-Lin Liu. Action recognition by dense trajectories. In Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition, CVPR '11, pages 3169–3176. IEEE Computer Society, 2011. 2
- [49] Jiang Wang, Yang Song, Thomas Leung, Chuck Rosenberg, Jingbin Wang, James Philbin, Bo Chen, and Ying Wu. Learning fine-grained image similarity with deep ranking. In *CVPR*, pages 1386–1393, 2014.
 4
- [50] X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. In *ICCV*, 2015. 4
- [51] Wanli Ouyang Cheng Ma Alan L. Yuille Xiaogang Wang Xiao Chu, Wei Yang. Multi-context attention for human pose estimation. In *CVPR*. IEEE, 2017. 6
- [52] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In Advances in neural information processing systems, pages 3320–3328, 2014. 6
- [53] Weiyu Zhang, Menglong Zhu, and Konstantinos G Derpanis. From actemes to action: A stronglysupervised representation for detailed action understanding. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2248–2255, 2013. 2

[54] Feng Zhou, Fernando De la Torre, and Jessica K Hodgins. Hierarchical aligned cluster analysis for temporal clustering of human motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(3):582–596, 2013. 2