Sympathy for the Details: Dense Trajectories and Hybrid Classification Architectures for Action Recognition

César Roberto de Souza[1,2(✉)], Adrien Gaidon[1], Eleonora Vig[3], and Antonio Manuel López[2]

[1] Computer Vision Group, Xerox Research Center Europe, Meylan, France
{cesar.desouza,adrien.gaidon}@xrce.xerox.com
[2] Centre de Visió per Computador, Universitat Autònoma de Barcelona, Bellaterra, Spain
antonio@cvc.uab.es
[3] German Aerospace Center, Wessling, Germany
eleonora.vig@dlr.de

Abstract. Action recognition in videos is a challenging task due to the complexity of the spatio-temporal patterns to model and the difficulty to acquire and learn on large quantities of video data. Deep learning, although a breakthrough for image classification and showing promise for videos, has still not clearly superseded action recognition methods using hand-crafted features, even when training on massive datasets. In this paper, we introduce hybrid video classification architectures based on carefully designed unsupervised representations of hand-crafted spatio-temporal features classified by supervised deep networks. As we show in our experiments on five popular benchmarks for action recognition, our hybrid model combines the best of both worlds: it is data efficient (trained on 150 to 10000 short clips) and yet improves significantly on the state of the art, including recent deep models trained on millions of manually labelled images and videos.

1 Introduction

Classifying human actions in real-world videos is an open research problem with many applications in multimedia, surveillance, and robotics [1]. Its complexity arises from the variability of imaging conditions, motion, appearance, context, and interactions with persons, objects, or the environment over different spatio-temporal extents. Current state-of-the-art algorithms for action recognition are based on statistical models learned from manually labeled videos. They belong to two main categories: models relying on features *hand-crafted* for action recognition (*e.g.,* [2–10]), or more recent end-to-end *deep architectures* (*e.g.,* [11–22]).

Electronic supplementary material The online version of this chapter (doi:10.1007/978-3-319-46478-7_43) contains supplementary material, which is available to authorized users.

B. Leibe et al. (Eds.): ECCV 2016, Part VII, LNCS 9911, pp. 697–716, 2016.
DOI: 10.1007/978-3-319-46478-7_43

These approaches have complementary strengths and weaknesses. Models based on hand-crafted features are data efficient, as they can easily incorporate structured prior knowledge (*e.g.,* the importance of motion boundaries along dense trajectories [2]), but their lack of flexibility may impede their robustness or modeling capacity. Deep models make fewer assumptions and are learned end-to-end from data (*e.g.,* using 3D-ConvNets [23]), but they rely on hand-crafted architectures and the acquisition of large manually labeled video datasets (*e.g.,* Sports-1M [12]), a costly and error-prone process that poses optimization, engineering, and infrastructure challenges.

Although deep learning for videos has recently made significant improvements (*e.g.,* [13,14,23]), models using hand-crafted features are the state of the art on many standard action recognition benchmarks (*e.g.,* [7,9,10]). These models are generally based on *improved Dense Trajectories* (iDT) [3,4] with Fisher Vector (FV) encoding [24,25]. Recent deep models for action recognition therefore combine their predictions with complementary ones from iDT-FV for better performance [23,26].

In this paper, we study an alternative strategy to *combine the best of both worlds via a single hybrid classification architecture* consisting in chaining sequentially the iDT hand-crafted features, the unsupervised FV representation, unsupervised or supervised dimensionality reduction, and a supervised deep network (*cf.* Fig. 1). This family of models was shown by Perronnin and Larlus [27] to perform on par with the deep convolutional network of Krizhevsky *et al.* [28]

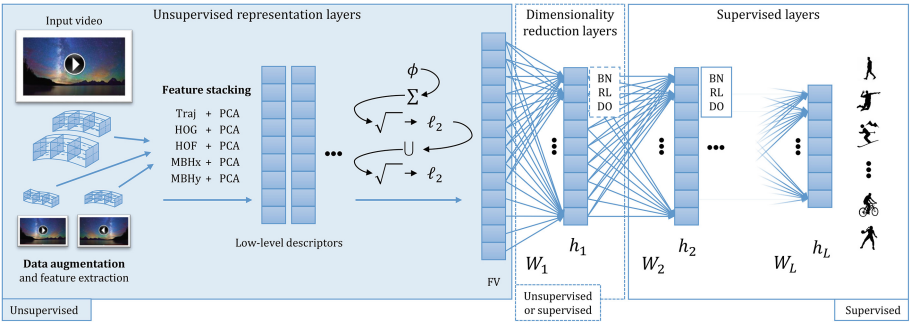

Fig. 1. Our hybrid unsupervised and supervised deep multi-layer architecture. Hand-crafted features are extracted along optical flow trajectories from original and generated videos. Those features are then normalized using RootSIFT [29], PCA-transformed, and augmented with their (x, y, t) coordinates, forming our low-level descriptors. The descriptors for each feature channel are then encoded (ϕ) as Fisher Vectors, separately aggregated (Σ) into a video-level representation, square-rooted, and ℓ_2-normalized. These representations are then concatenated ($\cup$) and renormalized. A dimensionality reduction layer is learned supervisedly or unsupervisedly. Supervised layers are followed by Batch-Normalization (BN) [30], ReLU (RL) non-linearities [31], and Dropout (DO) [32] during training. The last layer uses sigmoids (multi-label datasets) or softmax (multi-class datasets) non-linearities to produce action-label estimates.

for large scale image classification. We adapt this type of architecture differently for action recognition in videos with particular care for data efficiency.

Our **first contribution** consists in a careful design of the first *unsupervised* part of our hybrid architecture, which even with a simple SVM classifier is already on par with the state of the art. We experimentally observe that showing *sympathy for the details* (*e.g.,* spatio-temporal structure, normalization) and doing *data augmentation by feature stacking* (instead of duplicating training samples) are critical for performance, and that optimal design decisions generalize across datasets.

Our **second contribution** consists in a *data efficient hybrid architecture* combining unsupervised representation layers with a deep network of multiple fully connected layers. We show that *supervised mid-to-end learning of a dimensionality reduction layer together with non-linear classification layers* yields an excellent compromise between recognition accuracy, model complexity, and transferability of the model across datasets thanks to reduced risks of overfitting and modern optimization techniques.

The paper is organized as follows. Section 2 reviews the related works in action recognition. Section 3 presents the details of the first unsupervised part (based on iDT-FV) of our hybrid model, while Sect. 4 does so for the rest of the architecture and our learning algorithm. In Sect. 5 we report experimental conclusions from parametric studies and comparisons to the state of the art on five widely used action recognition datasets of different sizes. In particular, we show that our hybrid architecture improves significantly upon the current state of the art, including recent combinations of iDT-FV predictions with deep models trained on millions of images and videos.

2 Related Work

Existing action recognition approaches (*cf.* [1] for a recent survey) can be organized into four broad categories based on whether they involve *hand-crafted vs. deep-based* video features, and a *shallow vs. deep* classifier, as summarized in Table 1.

Table 1. Categorization of related recent action recognition methods

	Shallow classifier	Deep classifier
Hand-crafted features	[2–10]	[33], our method
Deep-based features	[23,26,34,35]	[11–22]

Hand-crafted features, shallow classifier. A significant part of the progress on action recognition is driven by the development of local hand-crafted spatio-temporal features encoded as bag-of-words representations classified by "shallow" classifiers such as SVMs [2–10]. Most successful approaches use *improved*

Dense Trajectories (iDT) [3] to aggregate local appearance and motion descriptors into a video-level representation through the Fisher Vector (FV) encoding [24,25]. Local descriptors such as HOG [36], HOF [37], and MBH [2] are extracted along dense point trajectories obtained from optical flow fields. There are several recent improvements to iDT, for instance, using motion compensation [5,6,38,39] and stacking of FVs to obtain a multi-layer encoding similar to mid-level representations [40]. To include global spatio-temporal location information, Wang *et al.* [5] compute FVs on a spatio-temporal pyramid (STP) [41] and use Spatial Fisher Vectors (SFV) [42]. Fernando *et al.* [10] model the global temporal evolution over the entire video using ranking machines learned on time-varying average FVs. Another recent improvement is the Multi-skIp Feature Stacking (MIFS) technique [7], which stacks features extracted at multiple frame-skips for better invariance to speed variations. An extensive study of the different steps of this general iDT pipeline and various feature fusion methods is provided in [8].

End-to-end learning: deep-based features, deep classifier. The seminal supervised deep learning approach of Krizhevsky *et al.* [28] has enabled impressive performance improvements on large scale image classification benchmarks, such as ImageNet [43], using Convolutional Neural Networks (CNN) [44]. Consequently, several approaches explored deep architectures for action recognition. While earlier works resorted to unsupervised learning of 3D spatio-temporal features [45], supervised end-to-end learning has recently gained popularity [11–22]. Karpathy *et al.* [12] studied several architectures and fusion schemes to extend 2D CNNs to the time domain. Although trained on the very large Sports-1M dataset, their 3D networks performed only marginally better than single-frame models. To overcome the difficulty of learning spatio-temporal features jointly, the Two-Stream architecture [13] is composed of two CNNs trained independently, one for appearance modeling on RGB input, and another for temporal modeling on stacked optical flow. Sun *et al.* [14] factorize 3D CNNs into learning 2D spatial kernels, followed by 1D temporal ones. Alternatively, other recent works use recurrent neural networks (RNN) in conjunction with CNNs to encode the temporal evolution of actions [16,17,19]. Overall, due to the difficulty of training 3D-CNNs and the need for vast amounts of training videos (*e.g.,* Sports-1M [12]), end-to-end methods report only marginal improvements over traditional baselines, and our experiments show that the iDT-FV often outperforms these approaches.

Deep-based features, shallow classifier. Several works [23,26,34,35] explore the encoding of general-purpose deep-learned features in combination with "shallow" classifiers, transferring ideas from the iDT-FV algorithm. Zha *et al.* [34] combine CNN features trained on ImageNet [43] with iDT features through a Kernel SVM. The TDD approach [26] extracts per-frame convolutional feature maps from two-stream CNN [13] and pools these over spatio-temporal cubes along extracted trajectories. Similar to [12], C3D [23] learns general-purpose features using a 3D-CNN, but the final action classifier is a linear SVM. Like

end-to-end deep models, these methods rely on large datasets to learn generic useful features, which in practice perform on par or worse than iDT.

Hybrid architectures: hand-crafted features, deep classifier. There is little work on using unsupervised encodings of hand-crafted local features in combination with a deep classifier. In early work, Baccouche *et al.* [33] learn temporal dynamics of traditional per-frame SIFT-BOW features using a RNN. The method, coupled with camera motion features, improves on BoW-SVM for a small set of soccer videos.

Our work lies in this category, as it combines the strengths of iDT-FV encodings and supervised deep multi-layer non-linear classifiers. Our method is inspired by the recently proposed hybrid image classification architecture of Perronnin and Larlus [27], who stack several unsupervised FV-based and supervised layers. Their hybrid architecture shows significant improvements over the standard FV pipeline, closing the gap on [28], which suggests there is still much to learn about FV-based methods.

Our work investigates this type of hybrid architectures, with several noticeable differences: (i) FV is on par with the current state of the art for action recognition, (ii) iDT features contain many different appearance and motion descriptors, which also results in more diverse and higher-dimensional FV, (iii) most action recognition training sets are small due to the cost of labeling and processing videos, so overfitting and data efficiency are major concerns. In this context, we adopt different techniques from modern hand-crafted and deep models, and perform a wide architecture and parameter study showing conclusions regarding many design choices specific to action recognition.

3 Fisher Vectors in Action: From Baseline to State of the Art

We first recall the iDT approach of Wang and Schmid [3], then describe the improvements that can be stacked together to transform this strong baseline into a state-of-the-art method for action recognition. In particular, we propose a data augmentation by feature stacking method motivated by MIFS [7] and data augmentation for deep models.

3.1 Improved Dense Trajectories

Local spatio-temporal features. The iDT approach used in many state-of-the-art action recognition algorithms (*e.g.,* [3–5,7,8,10,40]) consists in first extracting dense trajectory video features [2] that efficiently capture appearance, motion, and spatio-temporal statistics. Trajectory shape (Traj) [2], HOG [36], HOF [37], and MBH [2] descriptors are extracted along trajectories obtained by median filtering dense optical flow. We extract dense trajectories from videos in the same way as in [3], applying RootSIFT normalization [29] (ℓ_1 normalization followed by square-rooting) to all descriptors.

Unsupervised representation learning. Before classification, we combine the multiple trajectory descriptors in a single video-level representation by accumulating their Fisher Vector encodings (FV) [24,25], which was shown to be particularly effective for action recognition [5,46]. This high-dimensional representation is based on the gradient of a generative model, a Gaussian Mixture Model (GMM), learned in an *unsupervised* manner on a large set of trajectory descriptors in our case. We use $K = 256$ Gaussians as a good compromise between accuracy and efficiency [3–5]. We randomly sample 256,000 trajectories from the pool of training videos, irrespectively of their labels, to learn one GMM per descriptor channel using 10 iterations of EM. Before learning the GMMs, we apply PCA to the descriptors, reducing their dimensionality by a factor of two. After learning the GMMs, we extract FV encodings for all descriptors in each descriptor channel and combine these encodings into a per-channel, video-level representation using sum-pooling, *i.e.* by adding FVs together before normalization. In addition, we apply further post-processing and normalization steps, as discussed in the next subsection.

Supervised classification. When using a linear classification model, we use a linear SVM. As it is standard practice and in order to ensure comparability with previous works [3,7,26,47], we fix $C = 100$ unless stated otherwise and use *one-vs-rest* for multi-class and multi-label classification. This forms a strong baseline for action recognition, as shown by previous works [5,26] and confirmed in our experiments. We will now show how to make this baseline competitive with recent state-of-the-art methods.

3.2 Bag of Tricks for Bag-of-Words

Incorporating global spatio-temporal structure. Incorporating the spatio-temporal position of local features can improve the FV representation. We do not use spatio-temporal pyramids (STP) [41], as they significantly increase both the dimensionality of the representation and its variance [48]. Instead, we simply concatenate the PCA-transformed descriptors with their respective $(x, y, t) \in \mathbb{R}^3$ coordinates, as in [7,48]. We refer to this method as Spatio-Temporal Augmentation (STA). This approach is linked to the Spatial Fisher Vector (SFV) [42], a compact model related to soft-assign pyramids, in which the descriptor generative model is extended to explicitly accommodate the (x, y, t) coordinates of the local descriptors. When the SFV is created using Gaussian spatial models (*cf.* Eq. 18 in [42]), the model becomes equivalent to a GMM created from augmented descriptors (assuming diagonal covariance matrices).

Normalization. We apply signed-square-rooting followed by ℓ_2 normalization, then concatenate all descriptor-specific FVs and reapply this same normalization, following [7]. The double normalization re-applies square rooting, and is thus equivalent to using a smaller power normalization [25], which improves action recognition performance [49].

Multi-Skip Feature Stacking (MIFS). MIFS [7] improves the robustness of FV to videos of different lengths by increasing the pool of features with frame-skipped versions of the same video. Standard iDT features are extracted from those frame-skipped versions and stacked together before descriptor encoding, decreasing the expectation and variance of the condition number [7,50,51] of the extracted feature matrices. We will now see that the mechanics of this technique can be expanded to other transformations.

3.3 Data Augmentation by Feature Stacking (DAFS)

Data augmentation is an important part of deep learning [26,52,53], but it is rarely used with hand-crafted features and shallow classifiers, particularly for action recognition where duplicating training examples can vastly increase the computational cost. Common data augmentation techniques for images include the use of random horizontal flipping [26,52], random cropping [52], and even automatically determined transformations [54]. For video classification, [9,10] duplicate the training set by mirroring.

Instead, we propose to generalize MIFS to arbitrary transformations, an approach we call *Data Augmentation by Feature Stacking* (DAFS). First, we extract features from multiple transformations of an input video (frame-skipping, mirroring, etc.) that do not change its semantic category. Second, we obtain a large feature matrix by stacking the obtained spatio-temporal features prior to encoding. Third, we encode the feature matrix, pool the resulted encodings, and apply the aforementioned normalization steps along this pipeline to obtain a *single augmented video-level representation*.

This approach yields a representation that simplifies the learning problem, as it can improve the condition number of the feature matrix further than just MIFS thanks to leveraging data augmentation techniques traditionally used for deep learning. In contrast to data augmentation for deep approaches, however, we build a single more robust and useful representation instead of duplicating training examples. Note also that DAFS is particularly suited to FV-based representation of videos as pooling FV from a much larger set of features decreases one of the sources of variance for FV [55].

4 Hybrid Classification Architecture for Action Recognition

4.1 System Architecture

Our hybrid action recognition model combining FV with neural networks (cf. Fig. 1) starts with the previously described steps of our iDT-DAFS-FV pipeline, which can be seen as a set of *unsupervised layers*. The next part of our architecture consists of a set of L fully connected *supervised layers*, each comprising a dot-product followed by a non-linearity. Let h_0 denote the FV output from the last unsupervised layer in our hybrid architecture, h_{j-1} the input of layer

$j \in \{1, ..., L\}$, $h_j = g(W_j h_{j-1})$ its output, with W_j the corresponding parameter matrix to be learned. For intermediate hidden layers we use the Rectified Linear Unit (ReLU) non-linearity [31] for g. For the final output layer we use different non-linearity functions depending on the task. For multi-class classification over c classes, we use the softmax function $g(z_i) = \exp(z_i)/\sum_{k=1}^{c} exp(z_k)$. For multi-label tasks we consider the sigmoid function $g(z_i) = 1/(1 + exp(-z_i))$.

Connecting the last unsupervised layer to the first supervised layer can result in a much higher number of weights in this section than in all other layers of the architecture. Since this might be an issue for small datasets due to the higher risk of overfitting, we study the impact of different ways to learn the weights of this *dimensionality reduction layer*: either with unsupervised learning (*e.g.,* using PCA as in [27]), or by learning a low-dimensional projection end-to-end with the next layers of the architecture.

4.2 Learning

Unsupervised layers. Our unsupervised layers are learned as described in Sect. 3.1.

Supervised layers. We use the standard cross-entropy between the network output $\hat{y}$ and the ground-truth label vectors y as loss function. For multi-class classification problems, we minimize the categorical cross-entropy cost function over all n samples:

$$C_{cat}(y, \hat{y}) = -\sum_{i=1}^{n}\sum_{k=1}^{c} y_{ik} log(\hat{y}_{ik}), \tag{1}$$

whereas for multi-label problems we minimize the binary cross-entropy:

$$C_{bin}(y, \hat{y}) = -\sum_{i=1}^{n}\sum_{k=1}^{c} y_{ik} log(\hat{y}_{ik}) - (1 - y_{ik}) log(1 - \hat{y}_{ik}). \tag{2}$$

Optimization. For parameter optimization we use the recently introduced Adam algorithm [56]. Since Adam automatically computes individual adaptive learning rates for the different parameters of our model, this alleviates the need for fine-tuning of the learning rate with a costly grid-search or similar methods. Adam uses estimates of the first and second-order moments of the gradients in the update rule:

$$\theta_t \leftarrow \theta_{t-1} - \alpha \cdot \frac{m_t}{(1-\beta_1^t)\sqrt{\frac{v_t}{1-\beta_2^t}} + \epsilon} \quad \text{where} \quad \begin{array}{c} g_t \leftarrow \nabla_\theta f(\theta_{t-1}) \\ m_t \leftarrow \beta_1 \cdot m_{t-1} + (1-\beta_1) \cdot g_t \\ v_t \leftarrow \beta_2 \cdot v_{t-1} + (1-\beta_2) \cdot {g_t}^2 \end{array} \tag{3}$$

and where $f(\theta)$ is the function with parameters θ to be optimized, t is the index of the current iteration, $m_0 = 0$, $v_0 = 0$, and β_1^t and β_2^t denotes β_1 and β_2 to the power of t, respectively. We use the default values for its parameters $\alpha = 0.001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$ proposed in [56] and implemented in Keras [57].

Batch normalization and regularization. During learning, we use batch normalization (BN) [30] and dropout (DO) [32]. Each BN layer is placed immediately before the ReLU non-linearity and parametrized by two vectors γ and β learned alongside each fully-connected layer. Given a training set $X = \{x_1, x_2, ..., x_n\}$ of n training samples, the transformation learned by BN for each input vector $x \in X$ is given by:

$$BN(x;\gamma,\beta) = \gamma\frac{x-\mu_B}{\sqrt{\sigma_B^2+\epsilon}} + \beta \quad \text{where} \quad \mu_B \leftarrow \frac{1}{n}\sum_{i=1}^{n} x_i \ , \quad \sigma_B^2 \leftarrow \frac{1}{n}\sum_{i=1}^{n}(x_i-\mu_B)^2 \quad (4)$$

Together with DO, the operation performed by hidden layer j can now be expressed as $h_j = r \odot g(BN(W_j h_{j-1}; \gamma_j, \beta_j))$, where r is a vector of Bernoulli-distributed variables with probability p and $\odot$ denotes the element-wise product. We use the same DO rate p for all layers. The last output layer is not affected by this modification.

Dimensionality reduction layer. When unsupervised, we fix the weights of the dimensionality reduction layer from the projection matrices learned by PCA dimensionality reduction followed by whitening and ℓ_2 normalization [27]. When it is supervised, it is treated as the first fully-connected layer, to which we apply BN and DO as with the rest of the supervised layers.

Bagging. Since our first unsupervised layers can be fixed, we can train ensemble models and average their predictions very efficiently for bagging purposes [27,58,59] by caching the output of the unsupervised layers and reusing it in the subsequent models.

5 Experiments

We first describe the datasets used in our experiments, then provide a detailed analysis of the iDT-FV pipeline and our proposed improvements. Based on our observations, we then perform an ablative analysis of our proposed hybrid architecture. Finally, we study the transferability of our hybrid models, and compare to the state of the art.

5.1 Datasets

We use five publicly available and commonly used datasets for action recognition. We briefly describe these datasets and their evaluation protocols.

The **Hollywood2** [60] dataset contains 1,707 videos extracted from 69 Hollywood movies, distributed over 12 overlapping action classes. As one video can have multiple class labels, results are reported using the mean average precision (mAP).

The **HMDB-51** [61] dataset contains 6,849 videos distributed over 51 distinct action categories. Each class contains at least 101 videos and presents a high intra-class variability. The evaluation protocol is the average accuracy over three fixed splits [61].

The **UCF-101** [62] dataset contains 13,320 video clips distributed over 101 distinct classes. This is the same dataset used in the THUMOS'13 challenge [63]. The performance is again measured as the average accuracy on three fixed splits.

The **Olympics** [64] dataset contains 783 videos of athletes performing 16 different sport actions, with 50 sequences per class. Some actions include interactions with objects, such as *Throwing*, *Bowling*, and *Weightlifting*. Following [3,7], we report mAP over the train/test split released with the dataset.

The **High-Five** (a.k.a. TVHI) [65] dataset contains 300 videos from 23 different TV shows distributed over four different human interactions and a negative (no-interaction) class. As in [5,6,65,66], we report mAP for the positive classes (mAP+) using the train/test split provided by the dataset authors.

5.2 Detailed Study of Dense Trajectory Baselines for Action Recognition

Table 2 reports our results comparing the iDT baseline (Sect. 3.1), its improvements discussed in Sect. 3.2, and our proposed data augmentation strategy (Sect. 3.3).

Reproducibility. We first note that there are multiple differences in the iDT pipelines used across the literature. While [3] applies RootSIFT only on HOG, HOF, and MBH, in [7] this normalization is also applied to the Traj descriptor. While [3] includes Traj in their pipeline, [5] omits it. Additionally, person bounding boxes are used to ignore human motions when doing camera motion compensation in [5], but are not publicly available for all datasets. Therefore, we reimplemented the main baselines and compare our results to the officially published ones. As shown in Table 2, we successfully reproduce the original iDT results from [3,47], as well as the MIFS results of [7].

Improvements of iDT. Table 2 shows that double-normalization (DN) alone improves performance over iDT on most datasets without the help of STA. We show that STA gives comparable results to SFV+STP, as hypothesized in Sect. 3.2. Given that STA and DN are both beneficial for performance, we combine them with our own method.

Data Augmentation by Feature Stacking (DAFS). Although more sophisticated transformations can be used, we found that combining a limited number of simple transformations already allows to significantly improve the iDT-based methods in conjunction with the aforementioned improvements, as shown in the "iDT+STA+DAFS+DN" line of Table 2. In practice, we generate on-the-fly 7 different versions for each video, considering the possible combinations of frame-skipping up to level 3 and horizontal flipping.

Fine tuned and non-linear SVMs. Attempting to improve our best results, we also performed experiments both fine-tuning C and also using a Gaussian kernel while fine-tuning γ. However, we found that those two sets of experiments did not lead to significant improvements. As DAFS already brings results competitive with the current state of the art, we set those results with fixed

Table 2. Analysis of iDT baselines and several improvements

	UCF-101 %mAcc (s.d.)	HMDB-51 %mAcc (s.d.)	Hollywood2 %mAP	High-Five %mAP+ (s.d.)	Olympics %mAP
iDT [3]	84.8 [47][a, b]	57.2	64.3	-	91.1
Our reproduction	85.0 (1.32)[a, b]	57.0 (0.78)	64.2	67.7 (1.90)	88.6
iDT+SFV+STP [5]	85.7[a, b]	60.1[a]	66.8[a]	68.1[a, b]	90.4[a]
Our reproduction	85.4 (1.27)[a, b]	59.3 (0.80)[a]	67.1[a]	67.8 (3.78)[a, b]	88.3[a]
iDT+STA+DN [7]	87.3	62.1	67.0	-	89.8
Our reproduction	87.3 (0.96)[b]	61.7 (0.90)	66.8	70.4 (1.63)	90.7
iDT+STA+MIFS+DN [7]	89.1	65.1	68.0	-	91.4
Our reproduction	89.2 (1.03)[b]	65.4 (0.46)	67.1	70.3 (1.84)	91.1
iDT+DN	86.3 (0.95)[b]	59.1 (0.45)	65.7	67.5 (2.27)	89.5
iDT+STA	86.0 (1.14)[b]	60.3 (1.32)	66.8	70.4 (1.96)	88.2
iDT+STA+DAFS+DN	**90.6 (0.91)**[b]	**67.8 (0.22)**	**69.1**	**71.0 (2.46)**	**92.8**

iDT: Improved Dense Trajectories; SFV: Spatial Fisher Vector; STP: Spatio-Temporal Pyramids; STA: Spatio-Temporal Augmentation; MIFS: Multi-skIp Feature Stacking; DN: Double-Normalization; DAFS: Data Augmentation Feature Stacking; [a]without Trajectory descriptor; [b]without Human Detector.

C as our current shallow baseline (FV-SVM). We will now incorporate those techniques in the first unsupervised layers of our hybrid models.

5.3 Analysis of Hybrid Models

In this section, we start from hybrid architectures with unsupervised dimensionality reduction learned by PCA. For UCF-101 (the largest dataset) we initialize W_1 with $r = 4096$ dimensions, whereas for all other datasets we use the number of dimensions responsible for 99 % of the variance (yielding less dimensions than training samples).

We study the interactions between four parameters that can influence the performance of our hybrid models: the output dimension of the intermediate fully connected layers (*width*), the number of layers (*depth*), the dropout rate, and the mini-batch size of Adam (*batch*). We systematically evaluate all possible combinations and rank the architectures by the average relative improvement *w.r.t.* the best FV-SVM model. Training all 480 combinations for one split of UCF-101 can be accomplished in less than two days with a single Tesla K80 GPU. We report the top results in Table 3 and visualize all results using the parallel coordinates plot in Fig. 2. Our observations are as follows.

Table 3. Top-5 best performing hybrid architectures with consistent improvements

Depth	Width	Batch	UCF-101 %mAcc	HMDB-51 %mAcc	Hollywood2 %mAP	High-Five %mAP+	Olympics %mAP	Relative improv.
2	**4096**	**128**	91.6	68.1	72.6	**73.1**	**95.3**	**2.46 %**
2	4096	256	91.6	67.8	72.5	72.9	95.3	2.27 %
2	2048	128	91.5	68.0	72.7	72.7	94.8	2.21 %
2	2048	256	91.4	67.9	72.7	72.5	95.0	2.18 %
2	512	128	91.0	67.4	**73.0**	72.4	95.3	2.05 %
1	-	-	**91.9**	**68.5**	70.4	71.9	93.5	1.28 %
Best FV-SVM (*cf.* Table 2)			90.6	67.8	69.1	71.0	92.8	0.00 %

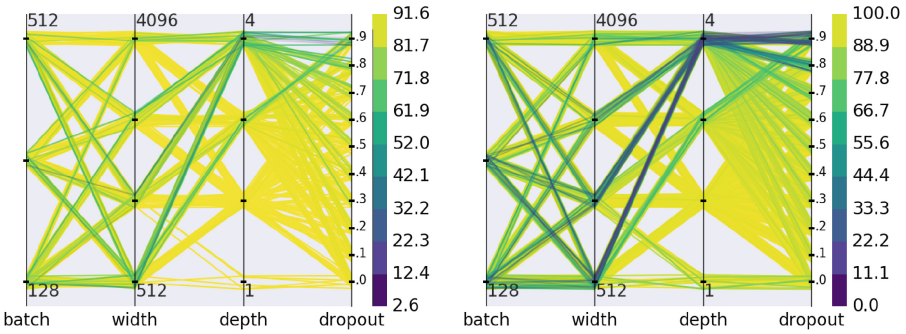

(a) UCF-101 (b) All datasets (normalized)

Fig. 2. Parallel coordinates plots showing the impact of multiple parameters. Each line represents one combination of parameters and color indicates performance of our hybrid architectures with unsupervised dimensionality reduction. Depth 2 correlates with high-performing architectures, whereas a small width and a large depth is suboptimal. (Color figure online)

Unsupervised dimensionality reduction. Performing dimensionality reduction using the weight matrix from PCA is beneficial for all datasets, and using this layer alone, achieves 1.28 % average improvement (Table 3, depth 1) upon our best SVM baseline.

Width. We consider networks with fully connected layers of size 512, 1024, 2048, and 4096. We find that a large width (4096) gives the best results in 4 of 5 datasets.

Depth. We consider hybrid architectures with depth between 1 and 4. Most well-performing models have depth 2 as shown in Fig. 2, but one layer is enough for the big datasets.

Dropout rate. We consider dropout rates from 0 to 0.9. We find dropout to be dependent of both architecture and dataset. A high dropout rate significantly impairs classification results when combined with a small width and a large depth.

Mini-batch size. We consider mini-batch sizes of 128, 256, and 512. We find lower batch sizes to bring best results, with 128 being the more consistent across all datasets. We observed that large batch sizes were detrimental to networks with a small width.

Best configuration with unsupervised dimensionality reduction. We find the following parameters to work the best: small batch sizes, a large width, moderate depth, and dataset-dependent dropout rates. The most consistent improvements across datasets are with a network with batch-size 128, width 4096, and depth 2.

Supervised dimensionality reduction. Our previous findings indicate that the dimensionality reduction layer can have a large influence on the overall

Table 4. Supervised dimensionality reduction hybrid architecture evaluation

Depth	Width	Batch	UCF-101 %mAcc (s.d.)	HMDB-51 %mAcc (s.d.)	Hollywood2 %mAP	High-Five %mAP+ (s.d.)	Olympics %mAP
1	1024	128	92.3 (0.77)	**69.4 (0.16)**	72.5	71.8 (1.37)	95.2
1	512	128	**92.3 (0.70)**	69.2 (0.09)	72.2	72.2 (1.14)	95.2
2	1024	128	91.9 (0.78)	68.8 (0.46)	71.8	72.0 (1.03)	94.8
2	512	128	92.1 (0.68)	69.1 (0.36)	70.8	71.9 (2.22)	94.2
Best unsup. (*cf.* Table 3)			91.9	68.5	**73.0**	**73.1**	**95.3**

classification results. Therefore, we investigate whether a *supervised* dimensionality reduction layer trained *mid-to-end* with the rest of the architecture could improve results further. Due to memory limitations imposed by the higher number of weights to be learned between our 116K-dimensional input FV representation and the intermediate fully-connected layers, we decrease the maximum network width to 1024. In spite of this limitation, our results in Table 4 show that much smaller hybrid architectures with supervised dimensionality reduction improve (on the larger UCF-101 and HMDB-51 datasets) or maintain (on the other smaller datasets) recognition performance.

Comparison to hybrid models for image recognition. Our experimental conclusions and optimal model differ from [27], both on unsupervised and supervised learning details (*e.g.,* dropout rate, batch size, learning algorithm), and in the usefulness of a supervised dimensionality reduction layer trained mid-to-end (not explored in [27]).

5.4 Transferability of Hybrid Models

In this section, we study whether the first layers of our architecture can be transferred across datasets. As a reference point, we use the first split of UCF-101 to create a base model and transfer elements from it to other datasets. We chose UCF-101 for the following reasons: it is the largest dataset, has the largest diversity in number of actions, and contains multiple categories of actions, including human-object interaction, human-human interaction, body-motion interaction, and practicing sports.

Unsupervised representation layers. We start by replacing the dataset-specific GMMs with the GMMs from the base model. Our results in the second row of Table 5 show that the transferred GMMs give similar performance to the ones using dataset-specific GMMs. This, therefore, greatly simplifies the task of learning a new model for a new dataset. We keep the transferred GMMs fixed in the next experiments.

Unsupervised dimensionality reduction layer. Instead of configuring the unsupervised dimensionality reduction layer with weights from the PCA learned on its own dataset, we configure it with the weights learned in UCF-101. Our results are in the third row of Table 5. This time we observe a different behavior: for Hollywood2 and HMDB-51, the best models were found without transfer,

Table 5. Transferability experiments involving unsupervised dimensionality reduction

Representation layers	Reduction layer	Supervised layers	HMDB-51 %mAcc (s.d.)	Hollywood2 %mAP	High-Five %mAP+ (s.d.)	Olympics %mAP
Own	Own	Own	68.0 (0.65)	**72.6**	73.1 (1.01)	95.3
UCF	Own	Own	**68.0 (0.40)**	72.4	73.7 (1.76)	94.2
UCF	UCF	Own	66.5 (0.88)	70.0	**76.3 (0.96)**	94.0
UCF	UCF	UCF	66.8 (0.36)	69.7	71.8 (0.12)	**96.0**

Table 6. Transferability experiments involving supervised dimensionality reduction

Representation layers	Supervised layers	HMDB-51 %mAcc (s.d.)	Hollywood2 %mAP	High-Five %mAP+ (s.d.)	Olympics %mAP
Own	Own	69.2 (0.09)	72.2	72.2 (1.14)	95.2
UCF	Own	69.4 (0.16)	**72.5**	71.8 (1.37)	95.2
UCF	UCF	**69.6 (0.36)**	72.2	**73.2 (1.89)**	**96.3**

whereas for Olympics it did not have any measurable impact. However, transferring PCA weights brings significant improvement in High-Five. One of the reasons for this improvement is the evidently smaller training set size of High-Five (150 samples) in contrast to other datasets. The fact that the improvement becomes less visible as the number of samples in each dataset increases (before eventually degrading performance) indicates there is a threshold below which transferring starts to be beneficial (around a few hundred training videos).

Supervised layers after unsupervised reduction. We also study the transferability of further layers in our architecture, after the unsupervised dimensionality reduction transfer. We take the base model learned in the first split of UCF-101, remove its last classification layer, re-insert a classification layer with the same number of classes as the target dataset, and fine-tune this new model in the target dataset, using an order of magnitude lower learning rate. The results can be seen in the last row of Table 5. The same behavior is observed for HMDB-51 and Hollywood2. However, we notice a decrease in performance for High-Five and a performance increase for Olympics. We attribute this to the presence of many sports-related classes in UCF-101.

Mid-to-end reduction and supervised layers. Finally, we study whether the architecture with supervised dimensionality reduction layer transfers across datasets, as we did for the unsupervised layers. We again replace the last classification layer from the corresponding model learned on the first split of UCF-101, and fine-tune the whole architecture on the target dataset. Our results in the second and third rows of Table 6 show that transferring this architecture brings improvements for Olympics and HMDB-51, but performs worse than transferring unsupervised layers only on High-Five.

5.5 Comparison to the State of the Art

In this section, we compare our best models found previously to the state of the art.

Best models. For UCF-101, the most effective model leverages its large training set using supervised dimensionality reduction (*cf.* Table 4). For HMDB-51 and Olympics, the best models result from transferring the supervised dimensionality reduction models from the related UCF-101 dataset (*cf.* Table 6). Due to its specificity, the best architecture for Hollywood2 is based on unsupervised dimensionality reduction learned on its own data (*cf.* Table 3), although there are similarly-performing end-to-end transferred models (*cf.* Table 6). For High-Five, the best model is obtained by transferring the unsupervised dimensionality reduction models from UCF-101 (*cf.* Table 5).

Bagging. As it is standard practice [27], we take the best models and perform bagging with 8 models initialized with distinct random initializations. This improves results by around one point on average, and our final results are in Table 7.

Discussion. In contrast to [27], our models outperform the state of the art, including methods trained on massive labeled datasets like ImageNet or Sports-1M, confirming both the excellent performance and the data efficiency of our approach. Table 8 illustrates some failure cases of our methods. Confusion matrices and precision-recall curves for all datasets are available in the supplementary material for fine-grained analysis.

Table 7. Comparison against the state of the art in action recognition

	Method	UCF-101 %mAcc (s.d.)	HMDB-51 %mAcc (s.d.)	Hollywood2 %mAP	High-Five %mAP+ (s.d.)	Olympics %mAP
HANDCRAFTED	iDT+FV [3]	84.8 [47]	57.2	64.3	-	91.1
	RCS [9]	-	-	73.6	71.1	-
	iDT+SFV+STP [5]	86.0	60.1	66.8	69.4	90.4
	iDT+MIFS [7]	89.1	65.1	68.0	-	91.4
	VideoDarwin [10]	-	61.6	69.6	-	-
	VideoDarwin+HF+iDT [10]	-	63.7	**73.7**	-	-
DEEP-BASED	2S-CNN [13]^{IN}	88.0	59.4	-	-	-
	2S-CNN+LSTM [17]^{IN}	88.6	-	-	-	-
	Objects+Motion(R*)[67]^{IN}	88.5	61.4	66.4	-	-
	Comp-LSTM [18]^{ID}	84.3	44.0	-	-	-
	C3D+SVM [23]^{S1M,ID}	85.2	-	-	-	-
	FSTCN [14]^{IN}	88.1	59.1	-	-	-
HYBRID	iDT+StackFV [40]	-	66.8	-	-	-
	TDD [26]^{IN}	90.3	63.2	-	-	-
	TDD+iDT [26]^{IN}	91.5	65.9	-	-	-
	CNN-hid6+iDT [34]^{S1M}	89.6	-	-	-	-
	C3D+iDT+SVM [23]^{S1M,ID}	90.4	-	-	-	-
	Best from state-of-the-art	91.5 [26]	66.8 [40]	**73.7** [10]	71.1 [9]	91.4 [7]
	Our best FV+SVM	90.6 (0.91)	67.8 (0.22)	69.1	71.0 (2.46)	92.8
	Our best hybrid	**92.5 (0.73)**	**70.4 (0.97)**	72.6	**76.7 (0.39)**	**96.7**

Methods are organized by category (cf. Table 1) and sorted in chronological order in each block. Our hybrid models improve upon the state of the art, and our handcrafted-shallow FV-SVM improves upon competing end-to-end architectures relying on external data sources (IN: uses ImageNet, S1M: uses Sports-1M, ID: uses private internal data).

Table 8. Top-5 most confused classes for our best FV-SVM and Hybrid models

Top-5 Confusions		#1	#2	#3	#4	#5
UCF-101	Hybrid	 P: ShavingBeard GT: BrushingTeeth	 P: FrisbeeCatch GT: LongJump	 P: ApplyLipstick GT: ShavingBeard	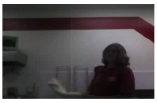 P: Nunchucks GT: PizzaTossing	 P: Rafting GT: Kayaking
	FV-SVM	 P: BreastStroke GT: FrontCrawl	 P: ShavingBeard GT: BrushingTeeth	 P: Rafting GT: Kayaking	 P: FrisbeeCatch GT: LongJump	 P: ApplyLipstick GT: ShavingBeard
HMDB-51	Hybrid	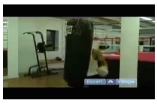 P: sword GT: punch	 P: flic_flac GT: cartwheel	 P: draw_sword GT: sword_exercise	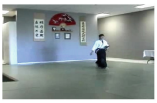 P: sword_exercise GT: draw_sword	 P: drink GT: eat
	FV-SVM	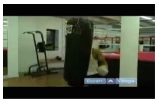 P: sword GT: punch	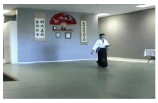 P: sword_exercise GT: draw_sword	 P: chew GT: smile	 P: throw GT: swing_baseball	 P: fencing GT: sword
High-Five	Hybrid	 P: negative GT: highFive	 P: negative GT: handShake	 P: hug GT: handShake	 P: hug GT: kiss	 P: handShake GT: highFive
	FV-SVM	 P: negative GT: highFive	 P: negative GT: handShake	 P: hug GT: kiss	 P: hug GT: handShake	 P: negative GT: kiss
Olympics	Hybrid	 P: long_jump GT: triple_jump	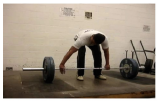 P: clean_and_jerk GT: snatch	 P: high_jump GT: vault	 P: discus_throw GT: hammer_throw	 P: vault GT: high_jump
	FV-SVM	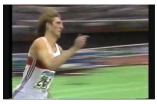 P: vault GT: high_jump	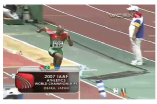 P: long_jump GT: triple_jump	 P: discus_throw GT: hammer_throw	 P: pole_vault GT: high_jump	 P: bowling GT: shot_put

6 Conclusion

We investigate hybrid architectures for action recognition, effectively combining hand-crafted spatio-temporal features, unsupervised representation learning based on the FV encoding, and deep neural networks. In addition to paying attention to important details like normalization and spatio-temporal structure, we integrate data augmentation at the feature level, end-to-end supervised dimensionality reduction, and modern optimization and regularization techniques. We perform an extensive experimental analysis on a variety of datasets, showing that our hybrid architecture yields data efficient, transferable models of small size that yet outperform much more complex deep architectures trained end-to-end on millions of images and videos. We believe our results open interesting new perspectives to design even more advanced hybrid models, *e.g.,* using recurrent neural networks, targeting better accuracy, data efficiency, and transferability.

Acknowledgements. Antonio M. Lopez is supported by the Spanish MICINN project TRA2014-57088-C2-1-R, and by the Secretaria d'Universitats i Recerca del Departament d'Economia i Coneixement de la Generalitat de Catalunya (2014-SGR-1506).

References

1. Vrigkas, M., Nikou, C., Kakadiaris, I.: A review of human activity recognition methods. Front. Robot. AI **2**, 1–28 (2015)
2. Wang, H., Kläser, A., Schmid, C., Liu, C.L.: Action recognition by dense trajectories. In: CVPR (2011)
3. Wang, H., Schmid, C.: Action recognition with improved trajectories. In: ICCV (2013)
4. Wang, H., Kläser, A., Schmid, C., Liu, C.L.: Dense trajectories and motion boundary descriptors for action recognition. IJCV **103**, 60–79 (2013)
5. Wang, H., Oneata, D., Verbeek, J., Schmid, C.: A robust and efficient video representation for action recognition. Int. J. Comput. Vision **119**(3), 219–238 (2015). doi:10.1007/s11263-015-0846-5
6. Gaidon, A., Harchaoui, Z., Schmid, C.: Activity representation with motion hierarchies. IJCV **107**, 219–238 (2014)
7. Lan, Z., Lin, M., Li, X., Hauptmann, A.G., Raj, B.: Beyond Gaussian pyramid: multi-skip feature stacking for action recognition. In: CVPR (2015)
8. Peng, X., Wang, L., Wang, X., Qiao, Y.: Bag of visual words and fusion methods for action recognition: comprehensive study and good practice. arXiv:1405.4506, May 2014
9. Hoai, M., Zisserman, A.: Improving human action recognition using score distribution and ranking. In: Cremers, D., Reid, I., Saito, H., Yang, M.-H. (eds.) ACCV 2014. LNCS, vol. 9007, pp. 3–20. Springer, Heidelberg (2015). doi:10.1007/978-3-319-16814-2_1
10. Fernando, B., Gavves, E., Oramas, M., Ghodrati, A., Tuytelaars, T., Leuven, K.U.: Modeling video evolution for action recognition. In: CVPR (2015)

11. Ji, S., Yang, M., Yu, K., Xu, W.: 3D convolutional neural networks for human action recognition. T-PAMI **35**, 221–31 (2013)
12. Karpathy, A., Leung, T.: Large-scale video classification with convolutional neural networks. In: CVPR (2014)
13. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NIPS (2014)
14. Sun, L., Jia, K., Yeung, D.Y., Shi, B.E.: Human action recognition using factorized spatio-temporal convolutional networks. In: ICCV (2015)
15. Wu, Z., Jiang, Y.g., Wang, X., Ye, H., Xue, X., Wang, J.: Fusing multi-stream deep networks for video classification. arXiv:1509.06086, September 2015
16. Donahue, J., Hendricks, L.A., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., Darrell, T.: Long-term recurrent convolutional networks for visual recognition and description. In: CVPR (2015)
17. Ng, J.Y.H., Hausknecht, M.J., Vijayanarasimhan, S., Vinyals, O., Monga, R., Toderici, G.: Beyond short snippets: deep networks for video classification. In: CVPR (2015)
18. Srivastava, N., Mansimov, E., Salakhutdinov, R.: Unsupervised learning of video representations using LSTMs. arXiv:1502.04681, March 2015
19. Ballas, N., Yao, L., Pal, C., Courville, A.C.: Delving deeper into convolutional networks for learning video representations. In: ICLR (2013)
20. Gan, C., Wang, N., Yang, Y., Yeung, D.Y., Hauptmann, A.G.: DevNet: A Deep Event Network for multimedia event detection and evidence recounting. In: CVPR (2015)
21. Wang, X., Farhadi, A., Gupta, A.: Actions $\sim$ transformations. In: CVPR (2015)
22. Feichtenhofer, C., Pinz, A., Zisserman, A.: Convolutional two-stream network fusion for video action recognition. In: CVPR (2016)
23. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3D convolutional networks. In: CVPR (2014)
24. Perronnin, F., Dance, C.: Fisher kernels on visual vocabularies for image categorization. In: CVPR (2007)
25. Perronnin, F., Sánchez, J., Mensink, T.: Improving the Fisher Kernel for large-scale image classification. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010. LNCS, vol. 6314, pp. 143–156. Springer, Heidelberg (2010). doi:10.1007/978-3-642-15561-1_11
26. Wang, L., Qiao, Y., Tang, X.: Action recognition with trajectory-pooled deep-convolutional descriptors. In: CVPR (2015)
27. Perronnin, F., Larlus, D.: Fisher vectors meet neural networks: a hybrid classification architecture. In: CVPR (2015)
28. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
29. Arandjelovi, R., Zisserman, A.: Three things everyone should know to improve object retrieval. In: CVPR (2012)
30. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: ICML (2015)
31. Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. In: ICML (2010)
32. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. **15**, 1929–1958 (2014)

33. Baccouche, M., Mamalet, F., Wolf, C., Garcia, C., Baskurt, A.: Action classification in soccer videos with long short-term memory recurrent neural networks. In: Diamantaras, K., Duch, W., Iliadis, L.S. (eds.) ICANN 2010. LNCS, vol. 6353, pp. 154–159. Springer, Heidelberg (2010). doi:10.1007/978-3-642-15822-3_20
34. Zha, S., Luisier, F., Andrews, W., Srivastava, N., Salakhutdinov, R.: Exploiting image-trained CNN architectures for unconstrained video classification. In: BMVC (2015)
35. Wu, Z., Wang, X., Jiang, Y., Ye, H., Xue, X.: Modeling spatial-temporal clues in a hybrid deep learning framework for video classification. In: ACM MM (2015)
36. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: CVPR (2005)
37. Dalal, N., Triggs, B., Schmid, C.: Human detection using oriented histograms of flow and appearance. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3952, pp. 428–441. Springer, Heidelberg (2006). doi:10.1007/11744047_33
38. Gaidon, A., Harchaoui, Z., Schmid, C.: Recognizing activities with cluster-trees of tracklets. In: BMVC (2012)
39. Jain, M., Jegou, H., Bouthemy, P.: Better exploiting motion for better action recognition. In: CVPR (2013)
40. Peng, X., Zou, C., Qiao, Y., Peng, Q.: Action recognition with stacked Fisher vectors. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 581–595. Springer, Heidelberg (2014). doi:10.1007/978-3-319-10602-1_38
41. Laptev, I., Marszalek, M., Schmid, C., Rozenfeld, B.: Learning realistic human actions from movies. In: CVPR (2008)
42. Krapac, J., Verbeek, J., Jurie, F.: Modeling spatial layout with Fisher vectors for image categorization. In: ICCV (2011)
43. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet large scale visual recognition challenge. IJCV **115**(3), 211–252 (2015)
44. LeCun, Y., Boser, B., Denker, J., Henderson, D., Howard, R., Hubbard, W., Jackel, L.: Handwritten digit recognition with a back-propagation network. In: NIPS (1989)
45. Le, Q.V., Zou, W.Y., Yeung, S.Y., Ng, A.Y.: Learning hierarchical invariant spatio-temporal features for action recognition with independent subspace analysis. In: CVPR (2011)
46. Chatfield, K., Lempitsky, V., Vedaldi, A., Zisserman, A.: The devil is in the details: an evaluation of recent feature encoding methods. In: BMVC (2011)
47. Wang, H., Schmid, C.: Lear-inria submission for the thumos workshop. In: ICCV workshop on action recognition with a large number of classes (2013)
48. Sanchez, J., Perronnin, F., De Campos, T.: Modeling the spatial layout of images beyond spatial pyramids. Pattern Recogn. Lett. **33**(16), 2216–2223 (2012)
49. Narayan, S., Ramakrishnan, K.R.: Hyper-Fisher Vectors for Action Recognition. arXiv:1509.08439 (2015)
50. Poggio, T., Smale, S.: The mathematics of learning: dealing with data. Not. Am. Math. Soc. **50**, 537–544 (2003)
51. Bousquet, O., Elisseeff, A.: Stability and generalization. J. Mach. Learn. Res. **2**, 499–526 (2002)
52. Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: delving deep into convolutional nets. In: BMVC (2014)

53. Konda, K., Bouthillier, X., Memisevic, R., Vincent, P.: Dropout as data augmentation. In: ICLR (2015)
54. Paulin, M., Harchaoui, Z., Perronnin, F., Paulin, M.: Transformation pursuit for image classification. In: CVPR (2014)
55. Boureau, Y.L., Ponce, J., LeCun, Y.: A theoretical analysis of feature pooling in visual recognition. In: ICML (2010)
56. Kingma, D., Ba, J.: Adam: a method for stochastic optimization. arXiv:1412.6980, December 2014
57. Chollet, F.: Keras: Deep learning library for Theano and TensorFlow (2015)
58. Maclin, R., Opitz, D.: An empirical evaluation of bagging and boosting. In: AAAI (1997)
59. Zhou, Z.H.H., Wu, J., Tang, W.: Ensembling neural networks: many could be better than all. Artif. Intell. **137**, 239–263 (2002)
60. Marszalek, M., Laptev, I., Schmid, C.: Actions in context. In: CVPR (2009)
61. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: HMDB: a large video database for human motion recognition. In: ICCV (2011)
62. Soomro, K., Zamir, A.R., Shah, M.: UCF101: A dataset of 101 human actions classes from videos in the wild. arXiv:1212.0402, December 2012
63. Jiang, Y.G., Liu, J., Zamir, R., a., Laptev, I., Piccardi, M., Shah, M., Sukthankar, R.: THUMOS Challenge: Action Recognition with a Large Number of Classes (2013)
64. Niebles, J.C., Chen, C.-W., Fei-Fei, L.: Modeling temporal structure of decomposable motion segments for activity classification. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010. LNCS, vol. 6312, pp. 392–405. Springer, Heidelberg (2010). doi:10.1007/978-3-642-15552-9_29
65. Patron-Perez, A., Marszalek, M., Zisserman, A., Reid, I.: High five: recognising human interactions in TV shows. In: BMVC (2010)
66. Patron-Perez, A., Marszalek, M., Reid, I., Zisserman, A.: Structured learning of human interaction in TV shows. T-PAMI **34**, 2441–2453 (2012)
67. Jain, M., Van Gemert, J., Snoek, C.G.M.: What do 15,000 object categories tell us about classifying and localizing actions? In: CVPR (2015)