

Learning Local Feature Aggregation Functions with Backpropagation

Angelos Katharopoulos*, Despoina Paschalidou*, Christos Diou, Anastasios Delopoulos
Multimedia Understanding Group

ECE Department, Aristotle University of Thessaloniki, Greece
{katharas, pdespoin}@auth.gr; diou@mug.ee.auth.gr; adelo@eng.auth.gr

Abstract—This paper introduces a family of local feature aggregation functions and a novel method to estimate their parameters, such that they generate optimal representations for classification (or any task that can be expressed as a cost function minimization problem). To achieve that, we compose the local feature aggregation function with the classifier cost function and we backpropagate the gradient of this cost function in order to update the local feature aggregation function parameters. Experiments on synthetic datasets indicate that our method discovers parameters that model the class-relevant information in addition to the local feature space. Further experiments on a variety of motion and visual descriptors, both on image and video datasets, show that our method outperforms other state-of-the-art local feature aggregation functions, such as Bag of Words, Fisher Vectors and VLAD, by a large margin.

I. INTRODUCTION

A typical image or video classification pipeline, which uses handcrafted features, consists of the following components: local feature extraction (e.g. Improved Dense Trajectories [15], SIFT [8]), local feature aggregation (e.g. Bag of Words [1], Fisher Vectors [10]) and classification of the final aggregated representation. This work focuses on the second component of the classification pipeline, namely the generation of discriminative global representations from the local image or video features.

The majority of existing local feature aggregation functions [1], [10], [13] rely on a visual codebook learned in an unsupervised manner. For instance, Bag of Words [1] quantizes every local feature according to a codebook, most commonly learned with K-Means, and represents the image as a histogram of codewords. Fisher Vectors [10], on the other hand, capture the average first and second order differences between the local feature descriptor and the centres of a GMM. Furthermore, the Kernel Codebook encoding [13] is analogous to the Bag of Words, with the only difference that it uses soft assignments, which are functions of the distances between the local features and the codewords.

There have been several attempts to improve the feature aggregation step by improving the codebook. For example, the authors of [4] propose a K-Means alternative that improves modelling of sparse regions of the local feature space. Other researchers focus on the indirect use of the class information in order to influence the codebook generation. For instance, in [6] Lazebnik et al. propose a technique for codebook learning

that aims to minimize the loss of the classification-relevant information. Finally, in [9] and [3] the authors make direct use of the class labels in order to improve the Bag of Words representation using a classifier.

In this paper, we define a family of local feature aggregation functions and we propose a method for the efficient estimation of their parameters in order to generate optimal representations for classification. In contrast to former research, our method:

- Can be used to estimate any type of parameters and not only codebooks.
- Can be used to create representations optimal for any task that can be expressed as a differentiable cost function minimization problem, not just classification.

To demonstrate these properties, we introduce two feature aggregation functions that outperform state-of-the-art local feature aggregation functions in terms of classification accuracy in various descriptors for both image and video datasets.

The rest of the paper is structured as follows. In Section II we introduce and explain the proposed method. Experimental results are reported in Section III, followed by conclusions in Section IV.

II. LEARNING LOCAL FEATURE AGGREGATION FUNCTIONS

Let $F = \{f_1, f_2, \dots, f_{N_F}\}$ be the set of N_F local descriptors extracted from an image or video. In order to derive a global representation for this feature set, we consider feature aggregation functions that can be expressed in the form of equation 1, where $T(\cdot; \Theta) : \mathbb{R}^D \mapsto \mathbb{R}^K$ is a differentiable function with respect to the parameters Θ .

$$R(F; \Theta) = \frac{1}{N_F} \sum_{n=1}^{N_F} T(f_n; \Theta) \quad (1)$$

By appropriately defining the $T(\cdot; \Theta)$ function, in the above formulation, we are able to express many local feature aggregation functions. For instance, the soft-assignment Bag of Words [13] can be expressed with the $T(\cdot; \Theta)$ function given in equation 2

$$T_{BOW}(f_n; C) = \frac{1}{\sum_{k=1}^K D(f_n, C_k)} \begin{bmatrix} D(f_n, C_1) \\ \vdots \\ D(f_n, C_K) \end{bmatrix} \quad (2)$$

* These two authors contributed equally

where

$$D(f_n, C_k) = \exp(-\gamma (f_n - C_k)^T (f_n - C_k)) \quad (3)$$

is a Gaussian-shaped kernel with Euclidean distance and $C \in \mathbb{R}^{D \times K}$ is the codebook.

In the following sections, we propose a generic method to estimate the parameters Θ^* of the local feature aggregation functions, such that they generate representations that are optimal for classification. To do that, we backpropagate the gradient of a classifier's cost function in order to update the parameters Θ using gradient descent.

A. Parameter estimation

Most approaches for parameter estimation of local feature aggregation functions do not take into consideration the subsequent usage of the global feature representation. For instance, in the case of the classification task, the extensively used K-Means and GMM methods, ignore the class labels of the feature vectors in the training set. In this work, we propose a supervised method for the parameter estimation of any local feature aggregation function that belongs in the family of functions of equation 1. Even though our method can be used for any task that can be expressed as a differentiable cost function minimization problem, in the rest of this paper we focus on the classification task. In particular, we estimate the values of the parameters Θ by minimizing the cost function $J(\cdot)$ of a classifier.

Let $J(x, y; W)$ be the cost function of a classifier with parameters W that aims to predict the class label y from a global feature vector x . Training a classifier is equivalent to finding the $W^* = \arg \min_W \frac{1}{N} \sum_{i=1}^N J(x^{(i)}, y^{(i)}; W)$, where $x^{(i)}$ and $y^{(i)}$ are the i -th training sample and its corresponding class label from a total of N samples. Instead of using traditional clustering methods, such as K-Means and GMM, to learn the parameters of the feature aggregation function, we compose $J(\cdot; W)$ with $R(\cdot; \Theta)$. This allows us to jointly learn a classifier and a feature aggregation function by solving the optimization problem of equation 4.

$$W^*, \Theta^* = \arg \min_{W, \Theta} \sum_{i=1}^N J\left(R(F^{(i)}; \Theta), y^{(i)}; W\right) \quad (4)$$

Due to the differentiability of $T(\cdot; \Theta)$, a straight-forward way to solve this optimization problem is to use Stochastic Gradient Descent (SGD). However, this optimization problem becomes computationally intensive in case of multimedia and especially for video datasets, due to the large number of local features F of each video (e.g. more than 20,000 local features in the case of Improved Dense Trajectories [15]). In order to address this problem, we approximate the gradient of $R(\cdot)$ with respect to the k -th parameter θ_k , of equation 5, by using a random sample of local features, S_F , instead of computing the gradient for every local feature.

$$\begin{aligned} \frac{\partial J}{\partial \theta_k} &= \frac{\partial J}{\partial R(F; \Theta)} \frac{\partial R(F; \Theta)}{\partial \theta_k} \\ &= \frac{\partial J}{\partial R(F; \Theta)} \frac{1}{N_F} \sum_{n=1}^{N_F} \frac{\partial T(f_n; \Theta)}{\partial \theta_k} \\ &\approx \frac{\partial J}{\partial R(F; \Theta)} \frac{1}{N_{S_F}} \sum_{n \in S_F} \frac{\partial T(f_n; \Theta)}{\partial \theta_k} \end{aligned} \quad (5)$$

Empirical results indicate that this approximation has similar effects to the stochastic gradient approximation of SGD, namely efficiency and robustness.

B. Aggregation functions

In this section, we make use of the previous analysis in order to create two local feature aggregation functions that outperform other state-of-the-art methods such as Bag of Words [1] and Fisher Vectors [10] on a variety of descriptors, as shown in the Experiments section III.

Firstly, we consider the representation $R_1(\cdot)$, which is a generalization of the soft-assignment Bag of Words and employs the encoding function $T_1(\cdot)$ of equation 6

$$T_1(f_n; C, \Sigma) = \frac{1}{Z(f_n, C, \Sigma)} \begin{bmatrix} D(f_n, C_1, \Sigma_1) \\ \vdots \\ D(f_n, C_K, \Sigma_K) \end{bmatrix} \quad (6)$$

where

$$D(f_n, C_k, \Sigma_k) = \exp(-\gamma (f_n - C_k)^T \Sigma_k^{-1} (f_n - C_k)) \quad (7)$$

and

$$Z(f_n, C, \Sigma) = \sum_{k=1}^K D(f_n, C_k, \Sigma_k) \quad (8)$$

involve the codebook C_k and the diagonal covariance matrix Σ_k used to compute the Mahalanobis distance between the n -th local feature and the k -th codeword.

On the other hand, we consider the representation $R_2(\cdot)$, produced by the encoding function $T_2(\cdot)$ of equation 9, which is exactly the soft-assignment Vector of Locally Aggregated Descriptors (VLAD) [2] and thus the dimensionality of the resulting representation is $D \times K$ because $f_n - C_k$ is a vector of size D .

$$T_2(f_n; C, \Sigma) = \frac{1}{Z(f_n, C, \Sigma)} \begin{bmatrix} D(f_n, C_1, \Sigma_1)(f_n - C_1) \\ \vdots \\ D(f_n, C_K, \Sigma_K)(f_n - C_K) \end{bmatrix} \quad (9)$$

In order to compute the optimal parameters C and Σ of the local feature aggregation functions, we optimize equation 4 using a Logistic Regression classifier with a cross-entropy loss according to equation 10. While linear classifiers are very efficient, non-linear classifiers tend to yield better classification results, especially in the case of Bag of Words [17]. Therefore, we decided to adopt an approximate feature map of χ^2 [14]

that is used in combination with $T_1(\cdot)$ and Logistic Regression to retain both the training efficiency of a linear classifier and the classification accuracy of a non-linear classifier.

$$J(x, y; W) = -\log \left(\frac{\exp(W_y^T x)}{\sum_{\hat{y}} \exp(W_{\hat{y}}^T x)} \right) \quad (10)$$

We could have used any classifier whose training is equivalent to minimizing a differentiable cost function, such as Neural Networks. Nevertheless, we use Logistic Regression and a χ^2 feature map in order to fairly compare our method to existing feature aggregation functions.

C. Training procedure

In Algorithm 1, we present the training procedure for the feature aggregation functions introduced in Section II-B. The training process consists of three main parts, the initialization step, the optimization step and the classifier fine-tuning step.

Regarding the initialization, we have experimented with three methods to initialize the codebook C and the covariance matrices Σ . In particular, we used:

- Random sampling from the set of local features to initialize the codebook and the identity matrix to initialize the covariance matrices.
- K-Means clustering to initialize the codebook and the identity matrix to initialize the covariance matrices.
- GMM clustering to initialize both the codebook and the covariance matrices

The proposed method can be used in combination to any of the aforementioned initializations. However, we empirically observe that when initialized with K-Means it results in a smoother parameter space, hence it is easier to choose a suitable value for the SGD learning rate. Finally, the reason for adding the classifier fine-tuning step emerged from the need to alleviate the effects of gradient noise, produced by the sampling of local features in equation 5.

III. EXPERIMENTS

This section presents an experimental evaluation of the proposed method on real and artificial datasets in order to assess its effectiveness and provide insights into the resulting feature aggregation functions. In particular, we have conducted experiments on the CIFAR-10 [5] image classification dataset and the UCF-11 (YouTube) Action dataset [7]. In case of CIFAR-10, we have extracted local features with a pre-trained deep convolutional neural network. Specifically, we have used the *conv3_3* layer from VGG-16 architecture [11], pre-trained on Imagenet, which results in 25 local features in \mathbb{R}^{256} for each image. In addition, in case of the video data, we have extracted Improved Dense Trajectories [15], after removing videos that have less than 15 frames, which results on an average of approximately 22,000 local features per video.

In Section III-A, we present a comparative evaluation of the discovered codewords in two synthetic datasets, in order to acquire a better understanding of the way our method chooses the codebook, compared to unsupervised methods.

Algorithm 1 Procedure to learn the parameters of a local feature aggregation function

```

procedure TRAINAGGFUN( $F, y$ )
  # Parameter initialization
  if initialize with K-Means then
     $C_0 \leftarrow KMeans(F)$ 
     $\Sigma_0 \leftarrow \mathbf{I}$ 
  else
     $C_0, \Sigma_0 \leftarrow GMM(F)$ 
  end if
   $W_0 \leftarrow \arg \min_W \sum_{i=1}^N J(R(F^{(i)}; C_0, \Sigma_0), y^{(i)}; W)$ 
  # Core training
   $t \leftarrow 0$ 
  repeat
     $i \sim \text{DiscreteUniform}(1, N)$ 
    Sample  $\hat{F}^{(i)}$  from  $F^{(i)}$ 
     $W_{t+1} \leftarrow \text{SGD}(\nabla_{W_t} J(R(\hat{F}^{(i)}; C_t, \Sigma_t), y^{(i)}; W_t))$ 
     $C_{t+1} \leftarrow \text{SGD}(\nabla_{C_t} J(R(\hat{F}^{(i)}; C_t, \Sigma_t), y^{(i)}; W_t))$ 
     $\Sigma_{t+1} \leftarrow \text{SGD}(\nabla_{\Sigma_t} J(R(\hat{F}^{(i)}; C_t, \Sigma_t), y^{(i)}; W_t))$ 
     $t \leftarrow t + 1$ 
  until  $t \geq$  specific number of mini-batches
  # Classifier fine tuning
   $C^* \leftarrow C_t$ 
   $\Sigma^* \leftarrow \Sigma_t$ 
   $W^* \leftarrow \arg \min_W \sum_{i=1}^N J(R(F^{(i)}; C^*, \Sigma^*), y^{(i)}; W)$ 
end procedure

```

Subsequently, in Section III-B, we present the classification accuracy of various representations on CIFAR-10, with respect to the training epochs, and compare it to the corresponding results using Bag of Words. Finally in Section III-C, we compare the proposed method on CIFAR-10 and UCF-11 with respect to the classification accuracy to Fisher Vectors, Bag of Words and VLAD on a variety of descriptors.

A. Synthetic dataset

Figure 1 compares the generated codebooks by K-Means, GMM and the proposed method on two artificial two-class datasets. In both cases, we generate and visualize 10 codewords, especially in the case of GMM we visualize additionally the covariance matrices. For our method, we use the $T_1(\cdot)$ feature aggregation function, from equation 6, to learn the codebook with the covariance matrices being fixed and equal to the identity matrix.

In contrast to K-Means and GMM, our method focuses on generating representations that can be separated by the classifier without necessarily retaining the structural information of the local features. It only suffices to observe Figure 1a to note that K-Means and GMM do not respect the circular class boundary, while our method focuses mainly on generating a linearly separable representation. In addition, owing to the fact that our method does not try to describe the local features it results in a more separable representation with a smaller codebook. For instance, it only requires a single codeword to successfully separate the concentric dataset of Figure 1a.




Fig. 1: Generated codebooks from synthetic data by K-Means, GMM and our method. The generated codewords are drawn with black crosses, while the dots are the local features from both classes (either blue or red).




Fig. 2: Classification accuracy on the test set with respect to the training epochs for various representation sizes on CIFAR-10 with features from VGG16 conv3_3

B. Training evolution

For this experiment, we generate codebooks using K-Means of sizes $\{64, 128, 256, 512, 1024, 2048\}$, which we subsequently use to create the corresponding Bag of Words representations. To classify the produced representations, we

train a linear SVM with a χ^2 feature map. Moreover, we use the $T_1(\cdot)$ feature aggregation function, of equation 6, with Logistic Regression, a χ^2 feature map and K-Means as an initialization method according to Algorithm 1. In order to select a value for the hyper-parameter γ of the $T_1(\cdot)$ function, we perform cross-validation.

By observing Figure 2, we conclude that the proposed method produces discriminative representations even with a small number of dimensions. In particular, it outperforms Bag of Words with 2048 dimensions by almost 4 percentage points with only 64 dimensions. Furthermore, we also notice that our method considerably improves the representation during the first epochs, thus we conclude that it can be used to fine-tune any differentiable feature aggregation function (e.g. Fisher Vectors) with little computational effort. Finally, we anticipate that increasing the number of training epochs will further increase the classification accuracy.

C. Classification results

In the current experiment, we assess the discriminativeness of the produced representations by evaluating their classification performance on a variety of descriptors and comparing it to several state-of-the-art feature aggregation methods. In the case of CIFAR-10, we use the provided train-test split while for UCF-11, we create three random 60/40 train-test splits and report both the mean classification accuracy and the

TABLE I: Comparison of $T_1(\cdot)$ and $T_2(\cdot)$ with Bag Of Words (BOW), VLAD and Fisher Vectors (FV)

	CIFAR-10	UCF-11	
	conv3_3	idt_hof	idt_traj
BOW-1024	69.12%	89.72% +/- 0.50	83.88% +/- 0.39
BOW-2048	70.50%	91.03% +/- 0.35	85.65% +/- 0.53
$T_1(\cdot)$ -1024 initial	76.85%	88.41% +/- 0.46	82.84% +/- 0.80
$T_1(\cdot)$ -1024	80.87%	92.23% +/- 0.37	86.90% +/- 0.63
$T_1(\cdot)$ -2048 initial	78.52%	90.24% +/- 0.24	83.78% +/- 0.90
$T_1(\cdot)$ -2048	81.12%	93.00% +/- 0.30	87.01% +/- 0.48
VLAD-64	-	90.25% +/- 0.33	78.71% +/- 0.94
FV-64	-	90.55% +/- 0.26	78.92% +/- 0.21
$T_2(\cdot)$ -64 initial	-	90.92% +/- 0.15	79.96% +/- 0.46
$T_2(\cdot)$ -64	-	91.08% +/- 0.26	83.82% +/- 0.34

standard error of the mean. Table I summarizes the results. The experimental setup for CIFAR-10 is analysed in Section III-B. Regarding UCF-11, we generate codebooks of sizes $\{1024, 2048\}$ using K-Means, both to create Bag of Words representations and to initialize the codebooks for the $T_1(\cdot)$ function. In addition, we train a GMM with 64 Gaussians to generate Fisher Vectors representations and again K-Means with 64 centroids to generate VLAD and initialize $T_2(\cdot)$.

For both datasets, we train an SVM with a χ^2 feature map for Bag of Words and $T_1(\cdot)$ and a linear SVM for the rest of the local feature aggregation functions in Table I. Moreover, in case of CIFAR-10, $T_1(\cdot)$ is trained for only 10 epochs, while for UCF-11, both $T_1(\cdot)$ and $T_2(\cdot)$ are trained for 30 epochs. In the conducted experiments, we have observed that both $T_1(\cdot)$ and $T_2(\cdot)$ are very sensitive with respect to the hyper-parameter γ , which must be carefully selected using a validation set or cross-validation. In particular, the reported results are generated using $\gamma = 70$ for UCF-11 “idt_traj”, $\gamma = 50$ for UCF-11 “idt_hof” and $\gamma = 5 \times 10^{-8}$ for CIFAR-10. The large differences in the range of γ make intuitive sense upon observing the distribution of the pairwise distances of the local features.

Furthermore, we additionally report the classification accuracy attained by $T_1(\cdot)$ and $T_2(\cdot)$, without learning the parameters using the proposed method; the results are reported in Table I as “initial”. This allows us to quantify the improvement in terms of classification accuracy achieved using the proposed method. In particular, we observe an average improvement of approximately 3.5 percentage points in all cases.

IV. CONCLUSIONS

We have introduced a new method to learn the parameters of a family of local feature aggregation functions through optimization, which can be used to learn any type of parameters and is not limited to codebooks. Furthermore, it can be used to generate an optimal representation for any task that can be expressed as a cost function minimization problem. In particular, in the conducted experiments, we have demonstrated the effectiveness of the proposed method in the classification task. We observed that the proposed local feature

aggregation functions outperform Bag of Words, Fisher Vectors and VLAD in a variety of descriptors on image and video data.

Our method opens up a multitude of new research directions. Initially, we could use the proposed method to learn extra parameters, such as γ , in order to further improve the generated representation. Moreover, it would be interesting to conduct experiments on other large-scale video classification datasets, such as UCF101 [12] and compare the performance of our method to state-of-the-art Neural Network architectures, such as the hybrid deep learning framework, as it was introduced in [16]. Finally, we can explore the use of the proposed method for the generation of optimal representations for other types of tasks, such as regression or ranking.

REFERENCES

- [1] G. Csurka, C. Dance, L. Fan, J. Willamowski, and C. Bray. Visual categorization with bags of keypoints. In *Workshop on statistical learning in computer vision, ECCV*, volume 1, pages 1–2. Prague, 2004.
- [2] H. Jégou, M. Douze, C. Schmid, and P. Pérez. Aggregating local descriptors into a compact image representation. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 3304–3311. IEEE, 2010.
- [3] M. Jiu, C. Wolf, C. Garcia, and A. Baskurt. Supervised learning and codebook optimization for bag-of-words models. *Cognitive Computation*, 4(4):409–419, 2012.
- [4] F. Jurie and B. Triggs. Creating efficient codebooks for visual recognition. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 1, pages 604–610. IEEE, 2005.
- [5] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009.
- [6] S. Lazebnik and M. Raginsky. Supervised learning of quantizer codebooks by information loss minimization. *IEEE transactions on pattern analysis and machine intelligence*, 31(7):1294–1309, 2009.
- [7] J. Liu, J. Luo, and M. Shah. Recognizing realistic actions from videos in the wild. In *Computer vision and pattern recognition, 2009. CVPR 2009. IEEE conference on*, pages 1996–2003. IEEE, 2009.
- [8] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.
- [9] F. Moosmann, B. Triggs, F. Jurie, et al. Fast discriminative visual codebooks using randomized clustering forests. In *NIPS*, volume 2, page 4, 2006.
- [10] F. Perronnin, J. Sánchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *European conference on computer vision*, pages 143–156. Springer, 2010.
- [11] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [12] K. Soomro, A. R. Zamir, and M. Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, abs/1212.0402, 2012.
- [13] J. C. Van Gemert, C. J. Veenman, A. W. Smeulders, and J.-M. Geusebroek. Visual word ambiguity. *IEEE transactions on pattern analysis and machine intelligence*, 32(7):1271–1283, 2010.
- [14] A. Vedaldi and A. Zisserman. Efficient additive kernels via explicit feature maps. *IEEE transactions on pattern analysis and machine intelligence*, 34(3):480–492, 2012.
- [15] H. Wang and C. Schmid. Action recognition with improved trajectories. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3551–3558, 2013.
- [16] Z. Wu, X. Wang, Y.-G. Jiang, H. Ye, and X. Xue. Modeling spatial-temporal clues in a hybrid deep learning framework for video classification. In *Proceedings of the 23rd ACM international conference on Multimedia*, pages 461–470. ACM, 2015.
- [17] J. Zhang, M. Marszałek, S. Lazebnik, and C. Schmid. Local features and kernels for classification of texture and object categories: A comprehensive study. *International journal of computer vision*, 73(2):213–238, 2007.