

Content-Based Filtering for Video Sharing Social Networks

Eduardo Valle¹, Sandra Avila², Fillipe de Souza²,
Marcelo Coelho^{2,3}, Arnaldo de A. Araújo²

¹ RECOD Lab — DCA / FEEC / UNICAMP, Campinas, SP, Brazil

² NPDI Lab — DCC / UFMG, Belo Horizonte, MG, Brazil

³ Preparatory School of Air Cadets — EPCAR, Barbacena, MG, Brazil

`dovalle@dca.fee.unicamp.br,`
`{sandra, fdms, mcoelho, arnaldo}@dcc.ufmg.br`

Abstract. *In this paper we compare the use of several features in the task of content filtering for video social networks, a very challenging task, not only because the unwanted content is related to very high-level semantic concepts (e.g., pornography, violence, etc.) but also because videos from social networks are extremely assorted, limiting the use of a priori information. We propose a simple method, able to combine diverse evidence, coming from different features and various video elements (entire video, shots, frames, keyframes, etc.). We evaluate our method in two social network applications, related to the detection of unwanted content — pornographic videos and violent videos. Using challenging test databases, we show that this simple scheme is able to obtain good results, provided that adequate features are chosen. Moreover, we establish the use of spatiotemporal local descriptors as critical to the success of the method in both applications.*

Resumo. *Neste trabalho, comparamos o uso de diferentes características na tarefa de filtragem de conteúdo para redes sociais de vídeo, uma tarefa muito desafiadora, não só porque o conteúdo indesejado está relacionado a conceitos semânticos de muito alto-nível (por exemplo, pornografia, violência, etc.), mas também porque os vídeos das redes sociais são extremamente variados, impedindo o uso de informação a priori. Propomos um método simples, capaz de combinar evidências diversa, provenientes de diferentes características e elementos de vídeo (vídeo inteiro, tomadas, quadros, quadros-chave, etc.) Avaliamos o nosso método em duas aplicações para redes sociais, relacionadas à detecção de conteúdos não desejados — vídeos pornográficos e vídeos violentos. Usando bases de dados de teste desafiadoras, mostramos que este esquema simples é capaz de obter bons resultados, desde que características adequadas sejam escolhidas. Além disso, mostramos que o uso de descritores espaço-temporais locais é crítico para o sucesso do nas duas aplicações.*

1. Introduction

Content-based classification and retrieval of visual documents by high-level semantic concepts has been an elusive goal pursued by the scientific community for the last 20 years. The persistent absence of a general solution attests the task difficulty, which is in great part brought by the much discussed “semantic gap” between the low-level representation of the data (pixels, frames, etc.) and the high-level concepts one wants to take into account.

We have, however, witnessed many important breakthroughs. In what concerns the description of visual documents, we have watched not only the inception and evolution of local features [22], but mainly the introduction of representations based on codebooks [5], which have allowed conciliating the discriminative power of the former with the generalization abilities required by high-level semantic tasks. Specifically for video, the introduction of “motion-aware” local features, which take into account the

dual nature of that media, at the same time spatial *and* temporal, has been an important achievement [6][13][14][15][16]. Meanwhile, the development of machine learning algorithms, like SVM [21], has created an effective framework for complex classification tasks.

In this paper, we are concerned with the detection of unwanted content on video sharing social networks — online communities built upon the production, sharing and watching of short video clips, which have been nourished by the popularization of broadband web access and the availability of cheap video acquisition devices. The crowds of users who employ the services of websites like *Dailymotion*, *MetaCafe* and *YouTube*, not only post and watch videos, but also share ratings, comments, “favorite” lists and other personal appreciation data.

The emergence of those networks has created a demand for specialized tools, including mechanisms to control abuses and terms-of-use violations. Indeed, the success of social networks has been inevitably accompanied by the emergence of users with non-collaborative behavior, which prevents them from operating evenly. Those behaviors include instigating the anger of other users (*trolling*, in the web jargon), diffusing materials of genre inappropriate for the target community (e.g., diffusing advertisement or pornography in inadequate channels), or manipulating illegitimately popularity ratings.

Non-collaborative behavior pollutes the communication channels with unrelated information, and prevents the virtual communities from reaching their original goals of discussion, learning and entertainment. It alienates legitimate users and depreciates the social network value as a whole [1].

In addition to the intricacies inherent to semantic classification, the challenges of content filtering are aggravated by the sheer amount of data social networks host and distribute. An automatic algorithm may be a strong ally to allow the detection of problematic content, but it is crucial that the number of false positives is kept low; otherwise the human agents will be overwhelmed.

In this paper, we address the posting of material considered inappropriate for the audience of the community, like pornographic or violent videos. Some hosts prohibit the posting of that material altogether, while others allow it, provided that it is especially flagged as “adult content”. Nevertheless, the content still ends up appearing where it is not welcome, because of either user ignorance of the rules, or full-fledged malice, when it is used to elicit revolted or shocked reactions from other users. Social networks face particular challenges, since content hosts and providers may face economic and, in some jurisdictions, even legal drawbacks (see §1 in [2]) if they not provide adequate means to protect their users from that kind of material.

We propose a simple scheme, inspired on voting algorithms, a popular technique which has been used in many tasks ranging from parameter estimation [34] to object detection [7] and video classification [35]. Using our scheme, we are able to combine several evidences (coming from classifiers using different features computed over different video elements) in order to obtain the decision on whether or not the video belonging to the unwanted class (e.g., pornography, violence, etc.). The idea is to ask several classifiers which label they attribute to the video. The “opinions” are counted, and the final decision is given by majority vote.

This paper presents two main contributions. First and foremost, the rigorous evaluation of several combinations of descriptors on two applicative scenarios, which has consistently indicated that a representation based on spatiotemporal bags of features is more discriminative than all alternative evidences.

In addition, we show that, with the right choice of features, a simple scheme, based on majority voting, is able to cope well with the difficult task of content filtering for social networks. We have evaluated our technique in very challenging datasets, conceived to represent the diversity of social networks, obtaining very promising results, with a good compromise between selectivity and specificity.

It is noteworthy that architecture proposed is very flexible, and can being easily adapted to any high-level concept the user might be interested in detecting.

2. Prior art

2.1. Video feature extraction

Semantic classification of visual documents has only become feasible after the emergence of effective feature extraction algorithms. Many of those may be applied to video, some being just still-image descriptors of individual frames, others being specially conceived to take into account the spatiotemporal nature of the moving image.

Though global image descriptors may be employed to characterize video frames, in the recent years a great deal of interest has been directed to local descriptors. Those are associated to different features of the image (regions, edges or small patches around points of interest) and have been shown to provide great robustness and discriminating power [17][18][22][23][24][25].

The most popular local descriptor, SIFT [7], is both a point of interest detector, based on differences of Gaussians and a local descriptor, based on the orientations of grayscale gradients. Using SIFT, visual content is represented by a set of scale and rotation invariant descriptors, which provides a characterization of local shapes. The generated descriptors allow for adequate levels of affine, viewpoint and illumination invariance.

Since color information is considered important for many tasks (e.g., nude detection), color extensions of SIFT have been proposed [9][19]. For example, a SIFT descriptor adapted to carry hue information (aptly named HueSIFT) had been proposed [9]. It provides color distinctiveness in addition to shape distinctiveness.

Intuition tells us that temporal information should be of prominent importance for recognition tasks in videos, for being likely to indicate interesting patterns of motion. Considering that, a few local features detectors and descriptors have been proposed, take into account the temporal nature of video [6][13][14][15][16]. For example, STIP [6] is designed as a differential operator, simultaneously considering extrema over spatial and temporal scales that correspond to particular patterns of events in specific locations. It extends the Harris corner detector [8] to the temporal domain, finding interest points as moving corner changes direction across a sequence — if the corner movement is constant, no interest point is detected. That allows detecting noteworthy “events” in the video sequence.

2.2. Codebooks of visual features

The discriminating power of local descriptors is extremely advantageous when matching objects in scenes, or retrieving specific target documents. However, when considering high-level semantic categories, it quickly becomes an obstacle, since the ability to generalize becomes then essential. A solution to that problem is to quantize the description spaces by using codebooks of local descriptors, in a technique sometimes named “visual dictionary”. The visual dictionary is nothing more than a representation that splits the descriptor space into multiple regions, usually by employing non-supervised learning techniques, like clustering. Each region becomes then a “visual word”.

The idea is that different regions of the description space will become associated to different semantic concepts, for example, parts of the human body, corners of furniture, vegetation, clear sky, clouds, features of buildings, etc. Yet, it is important to emphasize that this association is latent; there is no need to explicitly attribute meanings to the words. The technique has been employed successfully on several works for retrieval and classification of visual documents [1][5][11].

Once the codebook is obtained, description is greatly simplified, since it is no longer based on the exact value of descriptors, but only on their associated “words”. The condensed description may be, for example, a histogram or simply a set of the words the video contains. That has two advantages: the rougher description is better adapted to complex semantics; and the computational burthen is alleviated, since algorithms now operate on a single summarized description, instead of a myriad of individual local descriptors.

Building the dictionary requires the quantization of the description space, which can be obtained by a clustering algorithm. However, state-of-the-art clustering methods are seldom (if ever) conceived for the needs of visual dictionary construction: high-dimensional spaces, large datasets and a large number of clusters. The commonest choice found in the literature is a combination of aggressive sub-sampling of the dataset, dimensionality reduction using PCA (*Principal Component Analysis*), and clustering using a simple or hierarchical *k-means* algorithm with Euclidean distance. That typical choice however, may be considerable faulty on several grounds [20], and the design of good methods for visual dictionary creation is an active theme of investigation.

In addition to moderating the discriminating power of descriptors, the dictionaries allow adapting to visual documents techniques formerly available only to textual data. Among those borrowings, one of the most successful has been the technique of *bags of words* (which considers textual documents simply as sets of words, ignoring any inherent structure). The equivalent in the CBIR universe has been called *bags of visual words*, *bags of features* or *bags of visual features*, sometimes abbreviated as BoVF. It greatly simplifies document description, which becomes a histogram of the visual words it contains. The introduction of that technique had a huge impact on content-based retrieval and classification of visual documents [12].

The straightforward extension of BoVF to video uses individual frame images (or selected keyframes). That allows representing semantic concepts that are independent from motion. However, previous works in human annotation of video

databases [36], indicate that even for humans, many important concepts can only be adequately apprehended by taking into account the temporal aspects of video. Therefore, an interesting possibility is making codebooks of space-time local descriptors [6], which take into account the dynamic aspects of video. In that work we evaluate the performance of both static and “motion-aware” bags of features.

2.3. Video content filtering

The importance of pornography detection in visual documents is attested by the large literature on the subject. The vast majority of those works is based on the detection of human skin, and suffers from a high rate of false positives in situations of non-pornographic body exposure (like in sports). Some works use secondary criteria (like the shape of the detected skin areas, rejection of facial close-ups, etc.) to lower that rate. A comprehensive survey on skin-detection based methods may be found in [2].

Few methods have explored other possibilities. Bags of visual features (explained in the previous section) have been employed for many complex visual classification tasks, including pornography detection in images and videos [1][4][10]. Those works, however, have explored only bags of static features. Kim et al. [3] compare the effectiveness of several MPEG-7 features, but again, concentrate only on static features, ignoring those related to motion. Very few works have explored spatiotemporal features or other motion information for detection of pornography [37][38][39]. Jansohn [39] uses bags of static visual features and analysis of motion, including motion histograms, as separate evidences, but does not consider bags of spatiotemporal features

Violence detection has been addressed in works targeting applications as diverse as surveillance systems and movie rating. As expected, the application scope greatly affects how the problem is attacked: for example, in video-surveillance, movies are often black and white, noisy and silent [28][29]; in feature action movies, the soundtrack is often very indicative of the scene action, and so on [31][32] (to the point that some works are based solely on soundtrack evidences [30]).

In [28], a hierarchical approach for detection of violence in surveillance videos is proposed. Several actions involving two people are detected: fist fighting, kicking, hitting with objects, among others. The information of motion trajectories of the image structures in the scenes is obtained with the computation of acceleration measure vectors and their jerks. However, that method poses some limitations, failing for situations involving more than two people and when fighters fall down to the ground. Siebel and Maybank [29] developed a surveillance system for aiding human operators in monitoring undesirable events in a metro station. In [33], regions whose color indicated the presence of skin and blood are analyzed to detect aggressive actions in movies. Later, motion intensities of those regions of interest are computed, higher values indicating violence. We could not find any of approach dealing within the diversity of videos found in social networks, neither employing bags of visual features representations.

3. The proposed scheme

The proposed scheme is very simple and works by extracting elements from the video (shots, frames, keyframes, etc.), extracting features from those elements (global

features, bags of visual features based on local features, statistics, etc.) and training different classifiers for each type of feature used. In the classification phase, the classifier opinion is asked for each individual video element, and the final decision is reached by majority voting. The whole scheme is illustrated on Figure 1 and explained, in detail, below.

Pre-processing (video element and feature extraction) step:

1. The elements of each video are extracted (shots, frames, keyframes, etc.);
2. The features are extracted from the appropriate elements of the video. Those may be visual features, statistics, etc.

Training step:

1. A SVM classifier is created for each type of feature [21]. In our work, we have used a linear kernel (which in preliminary tests, has offered the best results);
2. Each classifier is trained with the corresponding features. Care is taken to balance the classes (positive and negative) so each is given roughly the same number of training samples at this step;

Classification step:

1. Each SVM classifier is asked about each single feature concerning all elements of the video related to that feature (i.e., if a feature is computed over keyframes, there will be a feature available for every keyframes, and the corresponding classifier will be asked once for each one of those features);
2. Every time it is enquired, a SVM classifier casts a vote: positive (the video is “unwanted”) or negative (the video is “ok”);
3. Those votes are counted for all classifiers on all features concerning the video. The majority label is given to the video.

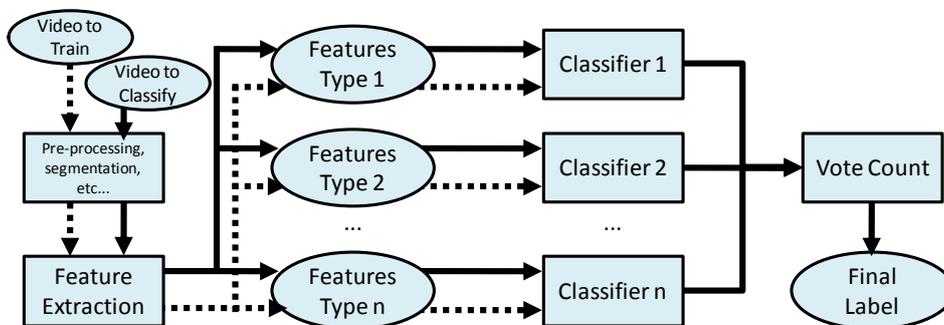


Figure 1: The proposed scheme for video classification. The data flow for training is represented by the dashed lines, while the data flow for classification is on continuous lines. Each classifier works on a different type of feature (e.g., color histogram, “bag” of local features, etc.) potentially computed over different video elements (frames, shots, etc.). The final label is obtained by majority voting over the opinion of all classifiers. That makes the scheme very robust.

4. First test application: pornography detection

Pornography is less straightforward to define than it may seem at first, since it is a high-level semantic category, not easily translatable in terms of simple visual characteristics. Though it certainly relates to nudity, pornography is a different concept: many activities

that involve a high degree of body exposure have nothing to do with it. That is why systems based on skin detection [2] often accuse false positives in contexts like beach shots or sports.

A commonly used definition is that pornography is the portrayal of *explicit* sexual matter with the *purpose* of eliciting arousal. That raises several challenges. First and foremost what threshold of explicitness must be crossed for the work to be considered pornographic? Some authors deal with that issue by further dividing the classes [1][3] but that not only falls short of providing a clear-cut definition, but also complicates the classification task. The matter of purpose is still more problematic, because it is not an objective property of the document.

4.1. Test database

We have opted to keep the evaluation conceptually simple, by assigning only two classes (porn and non-porn). On the other hand, we took great care to make them representative of the diversity found on social networks.

For the pornographic class, we have browsed social networks which only host that kind of material (solving, in a way, the matter of purpose) and sampled 400 videos as broadly as we could (Table I) — the database contains several genres of pornography and depicts actors of many ethnicities (Table II).

For the non-pornographic class we have browsed general-public social networks and selected two samples: 200 videos chosen at random (which we called “easy”) and 200 videos selected from textual search queries like “beach”, “wrestling”, “swimming”, which we knew would be particularly challenging for the detector (“difficult”).

Table I: A summary of the test database for pornography detection.

Class	Videos	Hours	Shots per Video
Porn	400	57	15.6
Non-Porn (“Easy”)	200	11.5	33.8
Non-Porn (“Difficult”)	200	8.5	17.5
All videos	800	77	20.6

Table II: Ethnic diversity on the pornographic videos.

Ethnicity	% of Videos
Asians	16 %
Blacks	14 %
Whites	46 %
Multi-ethnic	24 %

4.2. Experimental setup

For that application, the scheme has been parameterized as follows:

1. The two classes considered were porn (positive) and non-porn (negative). It is important to notice that the “easy” and “difficult” non-porn videos are considered in the same classification class. The differentiation is important only for the detailed analysis.

2. The video elements considered are video shots and the middle-frame of each shot. Video shots were obtained by an industry-standard segmentation software¹.
3. The following features computed for the frames:
 - Color Histogram: a normalized 64-bin RGB color histogram;
 - SIFT-BoVF: 5000-bin normalized BoVF using the SIFT descriptor [7];
 - HueSIFT-BoVF: a 5000-bin normalized BoVF using the HueSIFT descriptor [9];
4. The following feature was computed for the shots:
 - STIP-BoVF: a 5000-bin normalized BoVF using the STIP descriptor [6].

Obtaining a baseline to compare with our method was a major challenge since, in general, the numbers reported on the literature are not comparable from one work to another. Often, the databases are given only very cursory description, making next to impossible to make a fair assessment of the actual experimental conditions. Therefore, we have opted to compare ourselves to PornSeer Pro, an industry standard video pornography detection system, which is readily available for evaluation purposes². PornSeer Pro is based on the detection of specific features (like breast, genitals or the act of intercourse) on individual frames. It examines each individual frame of the video.

The experimental design was a classical 5-fold cross-validation, generating approximately 640 videos for training and 160 for testing on each fold.

4.3. Results

Figure 2 shows the performance, in the ROC space, of our detector using different combinations of features. It also shows the performance of the baseline method chosen, PornSeer. The graph reveals that several configurations of our detector are not significantly worse than PornSeer, and suggests that a few are significantly better. That latter statement, however, requires a more stringent statistical test, because we are comparing several configurations at once [26][27].

Therefore, we have performed an ANOVA test, using the 5 runs for all configurations shown in the graph. The model, using the configuration as a factor, was deemed as significant, with a p -value of less than 0.01 for both axes. That authorized us to perform pair-wise t-tests between the configurations, Table V shows the p -values obtained. We have highlighted values below 0.05 as significant.

The confusion matrix is another way to express the results shown in the ROC graphs. We showed the matrices for PornSeer (Table III) and for, arguably, the best configuration of our detector, using just the STIP descriptor (Table IV).

Table III. The average confusion matrix for PornSeer.

		Video was labeled as	
		Porn	Non-porn
Video was	Porn	65.1 %	34.9 %
	Non-porn	12.5 %	87.5 %

¹ <http://www.stoik.com/products/svc/>

² <http://www.yangsky.com/products/dshowseer/porndetection/PornSeePro.htm>

Table IV. The average confusion matrix for our scheme using STIP.

		Video was labeled as	
		Porn	Non-porn
Video was	Porn	91.3 %	8.7 %
	Non-porn	7.5 %	92.5 %

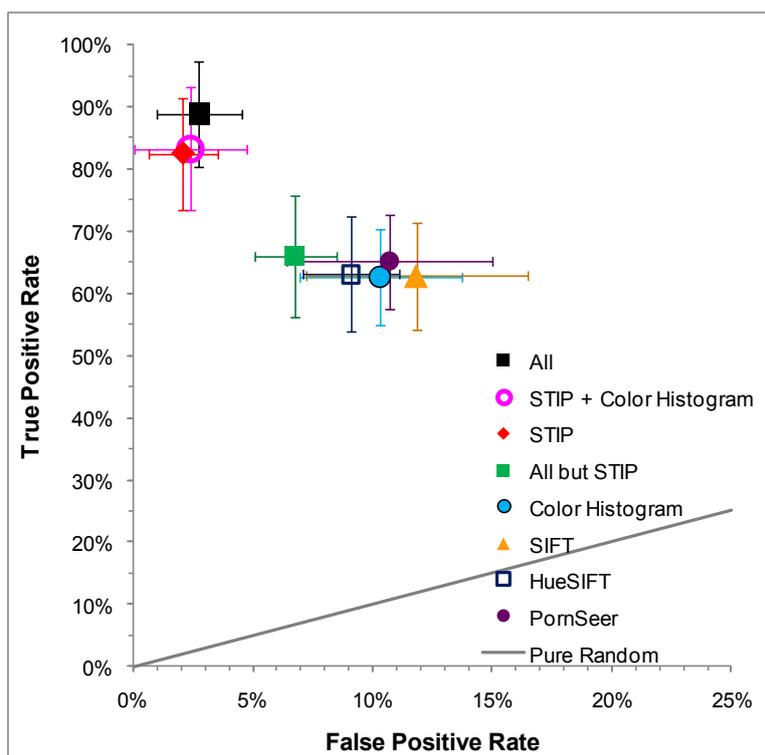


Figure 2: A few selected points of the above curves. The error bars are confidence intervals on the respective dimensions (for $\alpha = 0.05$). The graph shows that several configurations of our scheme are significantly better than the baseline (PornSeer).

Table V. P-values of pairwise t-test of detector configurations (and PornSeer), with significant differences in boldface. The pairwise tests were done after the ANOVA of all configurations was deemed significant (bottom row).

	True Positive Rates								False Positive Rates							
	1	2	3	4	5	6	7	8	1	2	3	4	5	6	7	8
1 – Color Hist.	—	.817	.459	.015	.049	.353	.041	.228	—	.169	.723	.020	.006	.763	.021	.289
2 – SIFT	.817	—	.333	.008	.030	.249	.024	.327	.169	—	.300	.001	<.001	.097	.001	.019
3 – Hue SIFT	.459	.333	—	.075	.202	.848	.174	.057	.723	.300	—	.009	.002	.513	.009	.161
4 – STIP	.015	.008	.075	—	.590	.108	.650	.001	.020	.001	.009	—	.589	.040	.993	.178
5 – STIP + Hist.	.049	.030	.202	.590	—	.275	.932	.003	.006	<.001	.002	.589	—	.012	.583	.064
6 – All but STIP	.353	.249	.848	.108	.275	—	.240	.038	.763	.097	.513	.040	.012	—	.041	.444
7 – All	.041	.024	.174	.650	.932	.240	—	.002	.021	.001	.009	.993	.583	.041	—	.181
8 – PornSeer	.228	.327	.057	.001	.003	.038	.002	—	.289	.019	.161	.178	.064	.444	.181	—
Model <i>p</i> -value				.0073								.0009				

4.4. Discussion

In its optimal configuration, our scheme is able to correctly identify 9 out of 10 of the pornographic clips, with few false positives. That is very important, since, as we have discussed, the cost of false alarms is high on the social network context, for it tends to overwhelm the human operators. The false positive rate attained may appear high at first, but it must be taken in the context of a very challenging dataset. Considering that half of the non-pornographic test videos were difficult cases, the rates are, actually, low.

It is instructive to study the cases where our method fails. The stubborn false positives correspond to very challenging non-pornographic videos: breastfeeding sequences, sequences of children being bathed, and beach scenes. The method succeeds for many videos with those subjects, but those particular ones have the additional difficulty of having very few shots (typically 1 or 2), giving no allowance for classification errors. PornSeer gave a wrong classification for all those clips.

The analysis of the most difficult false negatives revealed that the method has difficulty when the videos are of very poor quality (typical of amateur porn, often uploaded from webcams) or when the clip is only borderline pornographic, with few explicit elements. PornSeer also had difficulty with those clips, misclassifying many of them.

The study of Table V reveals interesting information. It shows that several configurations of our method have significantly better true positive rate than PornSeer, without increasing significantly the false positive rate.

More interestingly, it shows that STIP, the spatiotemporal descriptor, is critical in obtaining those good results. STIP used alone beats, in at least one of the axis, all configurations that do not use STIP; and it ties with all configurations that use STIP in combination with other descriptors. That suggests that not only spatiotemporal information is better for pornography detection in video, but also that combining it with other information (including color!) does not ameliorate the results.

Accumulation of evidences, however, seems to be useful when none of them is much compelling. Though none of the configurations using a single descriptors (other than STIP), is able to beat PornSeer, when the three “weak” descriptors (SIFT, HueSIFT and Color Histogram) are used together, they achieve significantly better detection rates.

5. Second test application: violence detection

Violence database suffers from the same problem of pornography detection: although it is a topic of great interest, with an abundant literature, the community lacks a shared violence dataset. In addition, existing works in do not describe the dataset used in enough details to allow a fair comparison. The matter is aggravated by the fact most works have more constrained applicative scope than ours: in the context of video-surveillance or feature movies, the data has more regularity to exploit than in the wild context of social networks, where there is less useful *a priori* information.

5.1. Test database

We have assembled a database containing 216 videos, 108 violent and 108 non-violent. The violent video clips come from very diverse contexts, including violent sports, street fights, civil unrest, etc. They also come from diverse sources: broadcasting, cell phones, video-surveillance cameras, etc.

5.2. Experimental setup

For this application, the scheme has been parameterized as follows:

1. The two classes considered were violent (positive) and non-violent (negative).
2. The video elements considered are video shots and the middle-frame of each shot. Video shots were obtained by an industry-standard segmentation software¹.
3. The following features computed for the frames:
 - SIFT-BoVF: 100-bin BoVF using the SIFT descriptor [7];
4. The following feature was computed for the shots:
 - STIP-BoVF: a 100-bin BoVF using the STIP descriptor [6].

The evaluation of the classification process was designed and conducted using the traditional 5-fold cross validation scheme, having approximately 160 videos for training and 40 for testing on each fold.

5.3. Results

The classification performances are presented in Table VI, for SIFT BoVF, and in Table VII for STIP BoVF. Though the results with SIFT are not bad (with over 80% accuracy for both classes) the results using STIP are impressive, scoring perfect results.

Table VI: Violent video classification using SIFT-BoVF.

		Video was labeled as	
		Violent	Non-violent
Video was	Violent	80.9%	19.1%
	Non-violent	5.0%	95.0%

Table VII: Violent video classification using STIP-BoVF.

		Video was labeled as	
		Violent	Non-violent
Video was	Violent	100%	0%
	Non-violent	0%	100%

5.4. Discussion

The analysis of the results indicates that local spatiotemporal features are decisive to distinguish between the violent and non-violent descriptors. A closer investigation at the SIFT features indicated that misclassification was, at least to some extent, due to cluttered backgrounds, low-quality frames and scenes of crowded people where the random poses make it extreme challenging to differentiate between violent and non-

violent situations. In those situations, the spatiotemporal descriptor (STIP) was able to provide additional motion information, allowing the classifier to reach the right conclusion. In fact, the spatiotemporal events typical of violent videos are so distinctive, that, though the classifier sometimes misses a shot or two, after majority voting is applied, no video is misclassified.

6. Conclusions

In both tasks, the spatiotemporal bags of features performed significantly better than all competing representations. That indicates that motion information is relevant for identifying the complex semantic categories present in the pornography and violence tasks. That result is not trivial: the state-of-art approaches on pornography detection, for example, are still heavily based upon (static) color and texture skin detection; violence detection approaches often use motion information without considering the advantages of semantic generalization that the codebook representation is able to provide.

When employing that recommended representation, the proposed scheme shows encouraging results in both tasks, even though the datasets employed were very challenging. A large fraction of unwanted videos is detected, without incurring in excessive false negatives.

In this article, we have evaluated different *visual* features as competing representations for the task of content filtering. Therefore, we have considered visual content as the main source of information. However, it is important to note that the proposed scheme disregards the media modality from where the features are extracted: textual, soundtrack and social interaction information could all be fed to the classifiers. In fact, we believe that, for extremely complex semantic tasks, a multimodal approach is needed to warrant the best performance possible.

Though our experiments indicate that incorporating information from other descriptors does not significantly improves the performance of the scheme using only STIP, we would like to explore non-trivial ways of incorporating that information. We were, for example, surprised by the fact the addition of color information did not improve the results (STIP is “color blind”, and, intuitively, color should be an important indicator of pornographic content). We would like to test if incorporating color directly inside the descriptor might improve the results.

The extraction of STIP features is currently extremely expensive, at around 1 *frame* per second! That severely limits the usefulness of the highly discriminant spatiotemporal descriptors for industrial applications. It is thus important to find ways to compute or to approximate those descriptors at reduced cost, especially for social networks/web-scale applications. The other steps of feature extraction are not expensive and, more importantly, scale well. The SVM classifier is relatively expensive to train, especially in terms of memory (the training set must be in RAM), but the trained model is cheaper to apply (the candidate element is compared only to the support vectors).

One interesting observation for both the pornography and violence applications is that label confidence is asymmetric on the training phase. When a training video is labeled as negative, we may be confident that none of its elements is positive. But a training video is labeled as positive may contain several negative shots, frames and keyframes where no violence or pornography is present (opening and closing credits,

“cut scenes”, etc.). Therefore, at least for the positive class, the training is only weakly supervised. Future versions of the scheme could take that information into account.

Acknowledgements

The authors thank CAPES, CNPq, FAPEMIG and FAPESP for the financial support that made this work possible. The local descriptors employed (SIFT, HueSIFT and STIP) were extracted with executable code provided by the authors of those methods.

References

- [1] T. Deselaers, L. Pimenidis and H. Ney. “Bag-of-Visual-Words Models for Adult Image Classification and Filtering”, In: *Proceedings of the Int. Conference on Pattern Recognition (ICPR'08)*, pp. 1-4, 2008.
- [2] W. Kelly, A. Donnellan, D. Molloy. “Screening for Objectionable Images: A Review of Skin Detection Techniques”, In: *Int. Machine Vision and Image Processing Conference (IMVIP'08)*, pp. 151-158, 2008.
- [3] W. Kim, S.J. Yoo, J-s. Kim, T.Y. Nam, and K. Yoon. “Detecting Adult Images Using Seven MPEG-7 Visual Descriptors”, In: *Web and Communication Technologies and Internet-Related Social Issues – HSI 2005*, PP. 336-339, 2005.
- [4] A. P. B. Lopes, S. E. F. de Avila, A. N. A. Peixoto, R. S. Oliveira, and A. de A. Araújo. “A Bag-of-Features Approach Based on Hue-SIFT Descriptor for Nude Detection”, In: *European Signal Processing Conference (EUSIPCO'09)*, pp. 1552-1556, 2009.
- [5] J. Sivic and A. Zisserman. “Video Google: A Text Retrieval Approach to Object Matching in Videos”, In: *Proceedings of the IEEE Int. Conference on Computer Vision (ICCV'03)*, pp. 1470-1477, 2003.
- [6] I. Laptev. “On Space-Time Interest Points”, In: *Int. Journal of Computer Vision (IJCV'05)*, vol 64, number 2/3, p.107-123, 2005.
- [7] D. G. Lowe. “Distinctive Image Features from Scale-Invariant Keypoints”, In: *Int. Journal of Computer Vision (IJCV'04)*, vol. 60, no. 2, pp. 91-110, 2004.
- [8] C. Harris and M. Stephens. “A Combined Corner and Edge Detector”, In: *Alvey Vision Conference*, pp. 147-152, 1998.
- [9] K. E. A. van de Sande, T. Gevers and C. G. M. Snoek. “Evaluating Color Descriptors for Object and Scene Recognition”, In: *IEEE Transactions on Pattern Analysis and Machine Intelligence (in press)*, 2010.
- [10] A. P. B. Lopes, S. E. F. Avila, A. N. A. Peixoto, R. S. Oliveira, M. M. Coelho and A. A. Araújo. “Nude Detection in Video Using Bag-of-Visual-Features”, In: *Proceedings of the 22th Brazilian Symposium on Computer Graphics and Image (SIBGRAPI)*, pp. 224-231, 2009.
- [11] Y. G. Jiang, C. W. Ngo and J. Yang. “Towards Optimal Bag-of-Features for Object Categorization and Semantic Video Retrieval”, In: *Proceedings of the 6th ACM Int. Conference on Image and Video Retrieval (CIVR '07)*, pp. 494-501, 2007.
- [12] J. Yang, Y.-G. Jiang, A. G. Hauptmann and C.-W. Ngo. “Evaluating Bag-of-Visual-Words Representations in Scene Classification”, In: *Proceedings of the Int. Workshop on Multimedia Information Retrieval (MIR'07)*, pp. 197-206, 2007.
- [13] P. Dollar, V. Rabaud, G. Cottrell, and S. Belongie, “Behavior Recognition via Sparse Spatio-Temporal Features”, In: *Proceedings of the 14th Int. Conference on Computer Communications and Networks (ICCCN'05)*, pp. 65-72, 2005.
- [14] H. Ning, Y. Hu and T. Huang. “Searching Human Behaviors Using Spatial-Temporal Words”, In: *Proceedings of the IEEE Int. Conference on Image Processing (ICIP'07)*, pp. 337-340, 2007.
- [15] J. C. Niebles, H. Wang, and L. Fei-Fei, “Unsupervised Learning of Human Action Categories Using Spatial-Temporal Words”, In: *Int. Journal of Computer Vision (IJCV'08)*, vol. 79, no. 3, pp. 299-318, 2008.
- [16] Y. Ke, R. Sukthankar, and M. Hebert, “Spatio-Temporal Shape and Flow Correlation for Action Recognition”, In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'07)*, pp. 1-8, 2007.

- [17] K. Mikolajczyk and C. Schmid. "An Affine Invariant Interest Point Detector", In: *Proceedings of the 7th European Conference on Computer Vision-Part I (ECCV'02)*, vol. 2350 of Lecture Notes in Computer Science, Springer Verlag, Berlin, Copenhagen, Denmark, pp. 128-142, 2002.
- [18] T. Tuytelaars and L. Van Gool. "Wide Baseline Stereo Matching Based on Local, Affinely Invariant Regions", In: *British Machine Vision Conference (BMVC'00)*, pp. 412-425, 2000.
- [19] A. E. Abdel-Hakim and A. A. Farag. "CSIFT: A SIFT Descriptor with Color Invariant Characteristics", In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, pp. 1978-1983, 2006.
- [20] F. Jurie and B. Triggs. "Creating Efficient Codebooks for Visual Recognition", In: *Proceedings of the IEEE Int. Conference on Computer Vision (ICCV'05)*, vol. 1, pp. 604-610, 2005.
- [21] C.-C. Chang and C.-J. Lin, "LIBSVM: A Library for Support Vector Machines", 2001, software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [22] T. Tuytelaars and K. Mikolajczyk. "Local Invariant Feature Detectors: A Survey", In: *Foundations and Trends in Computer Graphics and Vision*, vol. 3, no. 3, pp.177-280, 2008.
- [23] K. Mikolajczyk and C. Schmid. "Indexing Based on Scale Invariant Interest Points", In: *Proceedings of the IEEE Int. Conference on Computer Vision (ICCV'01)*, pp. 525-531, 2001.
- [24] C. Schuldt, I. Laptev and B. Caputo. "Recognizing Human Actions: A Local SVM Approach", In: *Proceedings of the Int. Conference on Pattern Recognition (ICPR'04)*, pp. III: 32-36, 2004.
- [25] K. Mikolajczyk and C. Schmid. "A Performance Evaluation of Local Descriptors", In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, n. 10, pp. 1615-1630, 2005.
- [26] J. Demšar. "Statistical Comparisons of Classifiers over Multiple Data Sets". In: *The Journal of Machine Learning Research*, vol. 7, 1-30, 2006.
- [27] S. Salzberg. "On Comparing Classifiers: Pitfalls to Avoid and a Recommended Approach". In: *Data Mining and Knowledge Discovery*, vol. 1, no. 3, 317-327, 1997.
- [28] A. Datta, M. Shah, and N. Da Vitoria Lobo, "Person-on-Person Violence Detection in Video Data", In: *Proceedings of the Int. Conference on Pattern Recognition (ICPR'02)*, vol. 1, p. 433-438, 2002.
- [29] N. T. Siebel and S. J. Maybank, "The Advisor Visual Surveillance System". In: *Proceedings of the ECCV 2004 Workshop Applications of Computer Vision (ACV'04)*, pp. 103-111, 2004.
- [30] T. Giannakopoulos, D. I. Kosmopoulos, A. Aristidou, and S. Theodoridis, "Violence Content Classification Using Audio Features", In: *Hellenic Artificial Intelligence Conference SETN-06, LNAI 3955*, pp. 502-507, 2006.
- [31] W. Zajdel, J. D. Krijnders, T. Andringa, and D. M. Gavrila, "CASSANDRA: Audio-Video Sensor Fusion for Aggression Detection", In: *IEEE Int. Conference on Advanced Video and Signal Based Surveillance (AVSS'07)*, pp. 200-205, 2007.
- [32] J. Lin and W. Wang, "Weakly-Supervised Violence Detection in Movies with Audio and Video based Co-Training", In: *Proceedings of the 10th Pacific Rim Conference on Multimedia (PCM'09)*, pp. 930-935, 2009.
- [33] C. Clarin, J. Dionisio, M. Echavez, and P. Naval, "Dove: Detection of Movie Violence Using Motion Intensity Analysis on Skin and Blood", In: *Proceedings of the 6th Philippine Computing Science Congress (PCSC'06)*, pp. 150-156, 2006.
- [34] D. Ballard. "Generalizing the Hough Transform to Detect Arbitrary Shapes". In: *Pattern Recognition*, vol. 13, n. 2, 111-122. Elsevier Inc., 1981.
- [35] W.-N. Lie, and C.-K. Su. "News Video Classification Based on Multi-Modal Information Fusion". In: *Proceedings of the IEEE Int. Conference on Image Processing (ICIP'06)*, 1213-1216, 2005.
- [36] L. Kennedy. "Revision of LSCOM Event/Activity Annotations", DTO Challenge Workshop on Large Scale Concept Ontology for Multimedia, Columbia University ADVENT Technical Report #221-2006-7, 2006.
- [37] T. Endeshaw, J. Garcia, and A. Jakobsson. "Classification of Indecent Video by Low Complexity Repetitive Motion Detection". In: *IEEE Applied Imagery Pattern Recognition Workshop*, pp. 1-7, 2008.
- [38] X. Tong, L. Duan, C. Xu, Q. Tian, L. Hanqing, J. Wang, and J. Jin. Periodicity Detection of Local Motion. In: *IEEE Int. Conference on Multimedia and Expo (ICME'05)*, pp. 650-653, 2005.
- [39] C. Jansohn, A. Ulges and T. M. Breuel. "Detecting Pornographic Video Content by Combining Image Features with Motion Information". In: *ACM Int. Conference on Multimedia*, pp. 601-604, 2009.