

# InFeRno - an Intelligent Framework for Recognizing Pornographic Web Pages

Sotiris Karavarsamis, Nikos Ntarmos and Konstantinos Blekas  
Department of Computer Science, University of Ioannina,  
PO Box1186, 45110 Ioannina, Greece,  
Email: {cs061205, ntarmos, kblekas}@cs.uoi.gr

## Abstract

In this work we present InFeRno, an intelligent web pornography elimination system, classifying web pages based solely on their visual content. The main characteristics of our system include: (i) a powerful vector space with a small but sufficient number of features that manage to improve the discriminative ability of the SVM classifier; (ii) an extra class (*bikini*) that strengthens the performance of the classifier; (iii) an overall classification scheme that achieves high accuracy at considerably lower runtime costs compared to current state-of-the-art systems; and (iv) a full-fledged implementation of the proposed system capable of being integrated with ICAP-aware web proxy cache servers.

**Keywords:** skin detection, contour extraction, SVM classifier

## 1 Introduction

Despite the usefulness and ease-of-access to a plethora of information scattered on the web, the Internet has become a hostile environment for unprotected people like children. Pornography is considered as sensitive information that is believed to be harmful for some groups of people. In the course of autonomously discriminating and blocking access to such content, Forsyth et al [1] was the first to implement a system comprising a figure grouper that

inferred the existence of nude human figures. However, a disadvantage of this method is its high processing time; typically it takes about 6 minutes to process a suspect image [2], that makes it impractical in real-world applications. Wang et al. [3] proposed another pornography elimination system called WIPE<sup>TM</sup>, which employed Daubechies wavelet analysis and extraction of invariant central moments.

Pornography filtering has also recently been applied to web searching upon user queries request. For instance, Rowley et al [4] have proposed a method for identifying nude images which has been part of Google<sup>TM</sup> safe search. Also, Hu et al. [5] proposed a system to recognize pornographic web pages by using both text and images information. It employs a combination of discrete and continuous text classifiers, a nearest-neighbor image classifier, as well as a method for fusing both kind of information. Similarly, the POESIA project [2] is an open-source system for blocking pornography, which applies a combination of image and harmful-symbol filtering and also text classification via NLP techniques.

In this work we present InFeRno, an intelligent web pornography elimination system, classifying web pages based solely on their visual content. The main characteristics of our system include: (i) a powerful vector space with a small but sufficient number of features that manage to improve the discriminative ability of the SVM classifier; (ii) an extra class (*bikini*) that strengthens the performance of the classifier; (iii) an overall classification scheme that achieves high accuracy at considerably lower runtime costs compared to current state-of-the-art systems; and (iv) a full-fledged implementation of the proposed system capable of being integrated with any ICAP<sup>1</sup>-aware web proxy cache server.

---

<sup>1</sup>ICAP stands for *Internet Content Adaptation Protocol* and is described in RFC 3507




Figure 1: InFeRno architecture

## 2 System architecture

The InFeRno core is implemented as an ICAP module, capable of communicating with any ICAP-enabled HTTP proxy software (see Fig 1). Client requests for web objects (pages or images) are forwarded to the proxy infrastructure which then delegates them to the InFeRno module. The latter first checks its local cache for an existing classification result. If one is found, it is returned to the client via the HTTP proxy; in the opposite case, InFeRno needs to fetch the requested object from the remote web server and classify it. For individual images, the result is returned to the client once the classification is over. In the case of web (html) pages, InFeRno further prefetches and classifies in parallel all (or a sample of) images referenced in the page code, then fuses the individual results to produce an overall classification for the web page. As image fetching and classification is done in parallel, the time required to classify a web page is very close to the time required to classify the “hardest” of its images (the overhead of the subsequent fusion step is in the order of a few milliseconds). Along with the caching of fetched objects and classification results, our system adds a slight overhead (of on average no more than 2”) in user-perceived web page loading times.



InFeRno can classify and filter content at various levels, at the discretion of its administrator: (i) at the web page level, it produces assessments for complete web pages (including all referenced images) and can deny access if the page is deemed pornographic (returning an appropriate error page to the

client), and (ii) at the image level, it can allow web pages to be rendered at the client, but deny access to individual images if they are deemed pornographic (returning a blurred out version of the image). The administrator can further tweak the acceptance thresholds for the amount of pornographic content being permitted per web page, thus allowing for an even more flexible configuration. The system core can also be integrated with other services, such as email scanner daemons, network firewalls, etc., providing image classification and porn filtering to a diverse array of end-user applications. Last, InFeRno can be configured either to not log client info so as to achieve a certain level of anonymization and privacy, or for full logging as an extra deterrent against illicit content access.

### 3 Classification system

**Skin detection.** InFeRno employs a simple rule-based skin color detection technique which uses deterministic rules imposing relations between the  $R$ ,  $G$  and  $B$  color channels [6]. As experiments have shown, this technique provides a satisfactory performance. However, it provides many false positives in the presence of either excessive illumination or of objects with skin-like color. For this reason, we have used an adaptive Gamma correction method, that eliminates the effect of this phenomenon and further improves the overall performance.

**Contour extraction.** In order to extract the contour of human nudes, we have followed a geometric approach, presented in [7], that uses a region splitting scheme. Initially, the image plane is partitioned into four equal quadrants. Subsequently, the corresponding skin and non-skin pixel intensity histograms in each quadrant are first found, and then two measurements are calculated: the *skin to non-skin ratio* and the *kurtosis*. Next, we decide to further split a region if both of the above measures exceed two predefined thresholds, estimated via extensive experimentation with diverse datasets.



(b) InFeRno vs. POESIA: response time in terms of image size

Figure 2: Operation and performance of InFeRno

The outer corners of the last constructed connected split-regions are regarded as control points which are finally used to establish the convex hull and localize the Region of Interest (ROI) in the image.

**Feature extraction and Classifier.** The next step is to extract a set of features regarding the size, orientation, and content of the contour that has been estimated previously. We have selected to calculate 15 such features, which are:

- Mean values and variances of the RGB color channels of all non-skin pixels.
- Ratio of the total skin and non-skin pixels delimited by the contour.
- 7 invariant spatial Hu moments of the surface.
- The angle of the diameter of the convex hull.

The above features are combined to form a feature vector space, which is used next for classifying images. We have selected the Support Vector Machine (SVM) as a classifier, using an one-against all scheme and simple linear

kernels, thus avoiding any free parameters. An innovation of our classification scheme is that we have considered one more class apart from pornographic and non-pornographic: the *bikini* class of people wearing bikinis, swimsuits, and so forth. Experiments have shown that introducing this third class manages to significantly not only increase the total performance of the classifier, but also strengthen its decision. Furthermore, the bikini class can optionally act as an extra level of nudity (a *softer* version of “porn”) recognized by the system. Fig 2(a) shows an example of the basic processing steps of our system.

For training the multi-class SVM, we have created a dataset of 5680 images (660 pornographic, 700 bikini and 4320 assorted benign images), that have been manually labeled. Using 10-fold cross validation we tested our classifier and received an accuracy of 98 %, 97 % and 98.8 % for our three classes, respectively. Last, Fig 2(b) depicts the time required to classify single images by both our system and the POESIA image classifier. Due to our system architecture and selected feature set, InFeRno achieves an equally high accuracy classification, but at a portion (more than a 4× speedup) of the time required by POESIA.

## References

- [1] M. Fleck, D. Forsyth, and C. Bregler, “Finding naked people,” *Computer Vision*, pp. 593–602, 1996.
- [2] POESIA project, “Poesia project.” <http://www.poesia-filter.org/>.
- [3] J. Z. Wang, J. Li, G. Wiederhold, and O. Firschein, “System for screening objectionable images,” *Computer Communications*, vol. 21, no. 15, pp. 1355–1360, 1998.

- [4] H. Rowley, Y. Jing, and S. Baluja, “Large scale image-based adult-content filtering,” in *Int. Conf. on Computer Vision Theory and Applications*, 2006.
- [5] W. Hu, O. Wu, Z. Chen, Z. Fu, and S. Maybank, “Recognition of pornographic web pages by classifying texts and images,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, pp. 1019–1034, 2007.
- [6] V. Vezhnevets, V. Sazonov, and A. Andreeva, “A survey on pixel-based skin color detection techniques,” in *Graphicon*, pp. 85–92, 2003.
- [7] J. Yang, Y. Shi, and M. Xiao, “Geometric feature-based skin image classification,” in *Int. Conf. on Adv. Intell. Comp. Theories and Applications*, pp. 1158–1169, 2007.