PPCensor: Architecture for real-time pornography detection in video streaming

Jackson Mallmann, Altair Olivo Santin, Eduardo Kugler Viegas, Roger Robson dos Santos, Jhonatan Geremias



PII:	S0167-739X(19)33107-3
DOI:	https://doi.org/10.1016/j.future.2020.06.017
Reference:	FUTURE 5683

To appear in: Future Generation Computer Systems

Received date : 22 November 2019 Revised date : 14 May 2020 Accepted date : 11 June 2020

Please cite this article as: J. Mallmann, A.O. Santin, E.K. Viegas et al., PPCensor: Architecture for real-time pornography detection in video streaming, *Future Generation Computer Systems* (2020), doi: https://doi.org/10.1016/j.future.2020.06.017.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2020 Published by Elsevier B.V.

# PPCensor: Architecture for Real-Time Pornography Detection in Video Streaming

Jackson Mallmann Federal Institute Catarinense, jackson.mallmann@ifc.edu.br Pontifical Catholic University of Paraná (PUCPR), jackson.mallmann@pucpr.edu.br

Brazil Eduardo Kugler Viegas Pontifical Catholic University of Paraná (PUCPR) eduardo.viegas@pucpr.br Brazil

Roger Robson dos Santos Pontifical Catholic University of Paraná (PUCPR) robson.roger@ppgia.pucpr.br Brazil

Altair Olivo Santin Pontifical Catholic University of Paraná (PUCPR) altair.santin@pucpr.br Brazil

> Jhonatan Geremias Pontifical Catholic University of Paraná (PUCPR) j.geremias@ppgia.pucpr.br Brazil

### ABSTRACT

Convolutional neural network (CNN) models are typically composed of several gigabytes of data, requiring dedicated hardware and significant processing capabilities for proper handling. In addition, video-detection tasks are typically performed offline, and each video frame is analyzed individually, meaning that the video's categorization (class assignment) as normal or pornographic is only complete after all the video frames have been evaluated. This paper proposes the Private Parts Censor (PPCensor), a CNN-based architecture for transparent and near real-time detection and obfuscation of pornographic video frame regions. Our contribution is two-fold. First, the proposed architecture is the first that addresses the detection of pornographic content as an object detection problem. The objective is to apply user-friendly content filtering such that an inevitable false positive will obfuscate only regions (objects) within the video frames instead of blocking the entire video. Second, the PPCensor architecture is deployed on dedicated hardware, and real-time detection is deployed using a video-oriented streaming proxy. If a pornographic video frame is identified in the video, the system can hide pornographic content (private parts) in real time without user interaction or additional processing on the user's device. Based on more than 50,000 objects labeled manually, the evaluation results show that the PPCensor is capable of detecting private parts in near real time for video streaming. Compared to cutting-edge CNN architectures for image classification, PPCensor achieved similar results, but operated in real time. In addition, when deployed on a desktop computer, PPCensor handled up to 35 simultaneous connections without the need for additional processing on the end-user device.

### **KEYWORDS**

Convolutional Neural Networks; Object-Oriented Private Parts Detection; Private Parts Object Dataset; Video-Oriented Streaming Proxy; Near Real-time Detection; Resource-Constrained Devices; Private Parts Censor.

### **1 INTRODUCTION**

The spread of pornographic videos on the internet, including those with explicitly pornographic content, has increased significantly in recent years [1]. For example, in 2018, one of the most visited pornographic sites received more than 4 million daily video uploads [2]. Exposure to this undesirable content, for example, by underage people without parental consent, can cause embarrassment or even psychological trauma [3]. Concerns regarding this kind of content are manifested by more than 71% of those responsible for underage children who are browsing the web [4].

Convolutional neural networks (CNNs) have, in general, demonstrated promising results in detecting pornographic video content [5–9]. For CNN analysis, a video is divided into a set of video frames (images), and each image is classified individually [10]. As each video consists of several frames and the typical frame rate is 23.97 frames per second (FPS) [11], current detection techniques cannot be used in real time, especially given that they require the entire set of video frames.

In real-world applications, videos are often streamed [12]. For instance, streaming pornographic websites (URL), such as LiveJasmin, Pornhub, XVideos, and xHamster, were included in the 100 most popular websites published by Alexa Internet in August 2019. In addition, because each video frame is classified individually, a single video can contain multiple frames labeled as pornographic, even if the video frame sequence does not contain inappropriate content for an underaged person.

A CNN-based detection approach may misclassify some video frames, and consequently, classify the video incorrectly [13]. Therefore, CNN-based video filtering can result in a high false positive (FP) rate, a situation in which several normal video frames are incorrectly classified as pornographic, as well as false negatives (FN), pornographic video frame deemed as normal. Moreover, CNN-based detection techniques often require a significant amount of memory and CPU/GPU capabilities [14]. One widely used CNN architecture can require up to 16 gigabytes of RAM while processing only 4.26 FPS on a high-end GPU [15].

CNN deployment on an end-user device has become challenging, considering that the available detection approaches ignore the limited

resources available for processing on end-user devices [16], including smartphones and tablets. As a result, current video filtering techniques based on the latest generation CNNs are unsuitable for deployment in such environments because their high computational demands typically require expensive and dedicated high-end hardware [13].

This paper presents PPCensor, an architecture for real-time pornographic object detection of video streaming frames. Our proposed technique is based on two main insights. First, to provide reliable real-time detection of private parts in video frames, our approach is the first that treats pornographic content as an object detection problem. The goal is to hide only the inappropriate areas within a set of frames instead of blocking access to the entire video sequence, as is often done in the literature. Consequently, PPCensor is less harsh than other approaches as the user retains access to the full video with potentially inappropriate objects being hidden. In addition, even if an FP occurs, the user can watch the affected frames with blurred (obfuscated) regions. Therefore, the user experience when using PPCensor is not significantly degraded owing to an FP. Second, to enable usage in production environments (real-world use), PPCensor is implemented as a network proxy server for video streams. The main advantage is the transparent, remote classification of the video streaming frames, i.e., running the CNN detection through a proxy instead of on the end-user device. PPCensor is deployed in dedicated hardware, and it evaluates the video content of all queried video URLs in real time. If pornographic video content is found, it may block access or hide inappropriate objects inside a frame without user intervention. PPCensor only requires the user to apply a proxy server to enable the device-agnostic filtering of inappropriate content.

The contributions of this study have been summarized below.

- We provide the first publicly available dataset with annotated private parts of a human body (objects), the Private Parts Object Dataset (PPO). This dataset was created by the manual inspection of pornographic video frames, resulting in 52,215 annotated objects. The dataset provides a benchmark for the research community that is building novel object-based pornographic detection approaches aimed at user-friendly detection and obfuscation of the private part.
- We present a novel technique that considers pornography detection as an object detection problem. The main advantage of the object detection approach is that we can hide (obfuscate) only the private parts (objects) within a video frame. Without using object detection, it would be necessary to remove or block entire video frames from the video stream to hide potential pornographic content, causing visual discontinuities in the media and degrading the user experience. Our proposed model achieves a low error rate (observed in the low FP and FN rates) and allows the implementation in production environments. Our results, which are the first real-time detection approach to be presented in the literature, highlight the feasibility of the proposed architecture that achieves accuracy rates similar to those of traditional offline techniques.
- We present a novel CNN-based detection architecture (PPCensor) suitable for mobile (resource-constrained) end-user devices. We implemented our prototype as a video streaming proxy, so that the CNN model is executed remotely. It does not require device modifications or additional processing on the device, and its working is transparent to the end-user.

The remainder of this paper is organized into the following sections. Section 2 presents work related to CNN-based detection techniques and object detection approaches. Section 3 evaluates current state-of-the-art pornographic video-detection techniques. Section 4 presents PPCensor, which is subsequently evaluated in Section 5. Finally, Section 6 summarizes our main conclusions.

### 2 RELATED WORK

Most video classification techniques split a video into several frames [10], and each frame is classified individually, applying an appropriate technique. These classifications are typically performed by three sequential modules: frame extractor, frame classification, and decision. The frame extractor module extracts the video frames to be classified. The frame classification module classifies each extracted video frame individually as normal or pornographic. Finally, the decision module classifies the video based on the individual frame classifications; in general, the video is classified according to the majority of the frame classifications.

### 2.1 Convolutional Neural Networks

In recent years, several detection techniques have been proposed for the frame classification task. In this context, CNN-based detection approaches have become cutting-edge procedures for image processing [13]. A CNN typically consists of three types of layers: convolutional, pooling, and fully connected [17]; its architecture varies in terms of organization and number of layers used.

The most important CNNs for image processing include the ResNet [18] and Inception [19] architectures. In both cases, it is necessary to initialize the configuration parameters, such as the number of epochs, learning rate, decay, and momentum [13]. To improve the detection accuracy and reduce the training time, several researchers have relied on pre-trained CNN architectures through a process known as transfer learning [20, 21]. For images, CNNs use the training weights obtained from the ImageNet dataset for transfer learning, a dataset with more than 1.2 million images divided into 1,000 classes [22].

The classification procedure begins with an input layer, which receives the raw image, i.e., a matrix of pixels, as input. The resulting data are passed to each subsequent layer until the output is reached. Finally, the output layer generates a probability vector for each classification label, i.e., classification probabilities are produced for normal or pornographic frames.

In the literature, a common approach to improving the accuracy of CNNs is to increase their network complexity. Several studies have considered increasing the number of layers, thereby increasing the demand for processing and memory size. It is common to rely on dedicated hardware equipped with the latest graphics processing units (GPUs). However, such approaches pose a challenge to resource-constrained devices.

Over the last several years, the detection of pornographic content in video frames has been widely studied. However, real-time detection is still in its infancy, especially for approaches intended for deployment on end-user devices. Most of the proposed approaches are based on skin tone or CNN [29].

A lightweight skin-tone-based approach, NuDetective [30], was proposed by Brazilian Federal Police experts. Their software uses the Ap-Apid algorithm [31] to automatically detect nudity in images according to the identified skin tones. NuDetective was developed as a forensic tool to help recover pornographic files from seized computers. In contrast, Nian et al. [32] detected pornographic images using a CNN architecture trained using transfer learning [20, 21] and a training dataset containing 48,600 images (13,300 pornographic and 35,300 normal) collected from the internet. The authors employed a medium-level representation and an adjustment of training data as their two main strategies. The main contribution of Nian's proposal is the application of the "fixed point algorithm" [33], which allows the detection of different image sizes.

The detection of pornographic content, in general, has been evaluated using publicly available datasets such as Pornography-2k [34], Sensitive [9], and NPDI [35]. Moustafa et al. [5] applied several CNN architectures, including AlexNet [22] and GoogleNet [36], for the detection of pornographic videos. In their proposal, videos were categorized as pornographic if the majority of their frames were classified as those containing inappropriate content, as evaluated based on the NPDI [35] dataset.

Another approach was proposed by Yahoo, namely, Yahoo Detector [6], which is a publicly available tool for finding not-safe-for-work images. Their technique applies a ResNet50 [18] CNN architecture trained using transfer learning [20, 21], based on ImageNet [22]. Although this technique yields reasonably accurate results, the used dataset was not shared, and therefore, cannot be compared with other approaches. In addition, Vitorino et al. [37] used a CNN architecture (GoogleNet [36]) trained using transfer learning and applied it to the publicly available Pornography-2k dataset [34]. Using this dataset, previously proposed techniques can be evaluated for accuracy comparisons. As an example, Yahoo Detector achieves an accuracy of only 88% when this dataset is used [37].

Perez et al. [7] evaluated their technique using the Pornography-2k [34] and NPDI [35] datasets. The authors proposed the use of static and motion-based features for CNN-based video classification. Their approach significantly improved the detection accuracy, yielding a classification accuracy of 97.9% for the NPDI [35] dataset. However, their approach cannot be implemented in real time because it requires all video frames for the extraction of motion-based features.

Wehrmann et al. [8] proposed a deep-learning architecture, called ACORDE, consisting of a CNN and recurring networks for pornographic content detection in video frames. Their experiments were performed using the NPDI [35] dataset and yielded reasonably accurate results, reducing the number of FPs and FNs of other approaches that utilized the same dataset.

In recent years, some authors have attempted to detect pornographic content in video frames in real time. For instance, Li et al. [50] proposed a system for the real-time detection of abnormal behaviors in live broadcasting platforms. The authors found abnormal rooms through several indicators, such as the degree of a scenery change, real-time comments, and threshold-based abnormal traffic detection. The authors were able to automatically identify abnormal rooms, but their technique ultimately relies on manual human inspection for detection.

Another approach was proposed by Wang et al. [51], which relies on multimodal features, such as video, motion, and sound, for realtime detection of pornographic content. In their study, the authors were able to provide a reasonable level of accuracy and a high detection throughput in broadcasting platforms, with a low FP rate. However, the authors relied on dedicated hardware for evaluation and thus did not address the execution of their technique on resource-constrained devices. Similarly, in a study conducted by Singh et al. [52], video scenes, i.e., sequences of frames, were detected, leveraging a long short-term memory autoencoder fed with a CNN video representation. Although the authors' proposal does not demand the entire video sequence for detection purposes, the deployment of their technique on resource-constrained devices was not addressed.

Although some proposals have enabled the real-time detection of pornographic content, the impacts of FPs on the user experience, as well as the deployment of such techniques on resource-constrained devices, remain an unaddressed challenge.

### 2.2 Object Detection Using CNN-Based Techniques

Traditional CNN-based video classification techniques classify a given input image as either belonging to or not belonging to a particular class (e.g., pornographic or normal). In specific applications, a given image may contain several regions of interest. An example application is a human detection system, in which a single input image may include several persons for identification. In such a case, the goal is the identification of several image regions as opposed to the classification of the entire image.

CNN-based approaches have also been used for object detection tasks in images [23]. The objective is to identify a bounding-box region (image region) that correctly includes the occurrence of the object in the image. Several approaches have been proposed to achieve such a goal, including faster R-CNN [25], single-shot detector (SSD) [26], and You Only Look Once detection [27]. This objective differs

from the goal of feature extractors that handles different sizes of input test images [24].

Similar to classification techniques, object detectors, such as faster R-CNNs [25] are also training-based. In Faster R-CNN, the features are first extracted from images and sent to a region proposal network (RPN). Second, the detection is performed for each location, computing the probability of the existence of the searched object. In contrast, the SSD architecture is executed in a single stage, considered fast and appropriate for resource-constrained hardware [24] such as smartphones and tablets.

To improve the detection accuracy, researchers frequently rely on a complex CNN architecture such as Faster R-CNN. As with imagebased classification, it is possible to train object detectors with transfer learning [20–21], which facilitates weight refinement using pretrained CNNs. In general, transfer learning in object detection is achieved using the COCO dataset [28], which consists of 330,000 images divided into 80 classes (objects).

In general, CNN-based object detectors are evaluated for their detection throughput and mean average precision (mAP). The mAP is computed according to the intersection over the union (IoU). In the literature, in general, it is applied at a threshold of 0.5 for each object bounding box. The IoU establishes the region of the identified bounding box that intersects with the target object region [23].

In recent years, several CNN architectures have been proposed for object detection tasks and utilized in several different fields. For instance, used for the recognition of vehicles [38–40], people [41], water surface objects [42], and have even been used for crime scene analysis [43].

Braun et al. [53] applied state-of-the-art object detection CNNs, including Faster R-CNN and SSD, for human recognition. In their study, the authors were able to reach a reasonable level of accuracy when detecting a person in an image. In contrast, for the identification of individual parts, Yang et al. [54] applied an SSD architecture for hand identification, reaching a reasonable accuracy rate. Another approach was proposed by Guo et al. [55] for the identification of facial regions in images through an object-detection CNN architecture.

The applicability of CNN-based object detection techniques for the identification of private parts for pornographic detection remains an open challenge. For object detection in pornography, Ou et al. [9] attempted to classify images/videos by utilizing a deep multi-context network (DMCNet) [44], which is a hierarchical method that applies the feature output of a CNN model to train the RPN in Faster R-CNN object detection. For the purpose of comparison, the authors generated a Sensitive dataset [9] and compared it with other datasets. Although the authors applied an object-detection CNN architecture, they did not detect any private parts.

Several approaches have been proposed for the detection of pornographic content in video frames. However, their implementation in real-world applications remains an open issue because to increase the level of accuracy, researchers typically extract features over several frames [5, 6] or significantly increase the CNN complexity [7–9]. Consequently, real-time detection is still very limited in the literature.

A pornography detection approach based on object detection techniques can significantly improve the user experience because an FP only causes an area of the video to be obfuscated instead of the entire set of video frames being restricted. There are only a few studies that have used object detection to locate private parts or to detect pornographic images.

Our proposal, PPCensor, presents two main benefits when compared to other state-of-the-art techniques. First, unlike traditional pornographic content detection, this is the first work to tackle pornographic detection as an object detection problem. As our contribution to this field, private human parts can be individually identified and properly obfuscated inside a video frame. Consequently, when an FP occurs, only a small section of the video frame is obfuscated; the entire video is not blocked, as in the case of current techniques. Second, unlike related work, we consider the resource-restricted nature of the end user's device. PPCensor runs as a video streaming proxy and is applied transparently, not generating additional processing for the end user's device.

### **3 PROBLEM STATEMENT**

In this section, we evaluate the classification performance of publicly available detection approaches for pornographic content. More specifically, we first describe the used pornographic datasets and then evaluate the accuracy and detection throughput of the NuDetective [30] and Yahoo Detector [6] techniques.

### 3.1 Datasets

The methods used for detecting pornography are typically evaluated against a single dataset. Thus, the accuracy obtained is not demonstrated in production (real-world) environments. In addition, the use of a single dataset can introduce overfitting in detection techniques because a single dataset is prone to several different types of specificities such as the environment in which it was acquired, the extracted quality and encoding, the video length, and other media characteristics.

We evaluated the detection of pornographic content in video frames using two datasets. Specifically, we considered the Pornography-2k [34] and UCF101 [45] datasets. The prior dataset is widely used as a reference for pornographic detection schemes and consists of 2,000 videos, of which 1,000 are related to pornography, and the other 1,000 are normal videos. In contrast, the latter dataset consists of 13,320 normal videos. The objective is to obtain a realistic detection reference for the FP and FN rates. The FP rate indicates the proportion of normal video frames incorrectly classified as pornographic, whereas the FN rate indicates the proportion of pornographic video frames incorrectly classified as normal. Both datasets include individuals of different ethnicities involved in different activities and with varying levels of video quality.

The video frames of the dataset were extracted at an interval of 10 frames, i.e., a frame was extracted for further analysis approximately every 0.3 seconds. In total, the Pornography-2k dataset consisted of 1,376,041 frames, of which 389,808 frames are normal, and 986,233 frames are pornographic, whereas the UCF101 [45] dataset consisted of videos with 240,810 normal frames.

#### 3.2 Reliability of Pornography Detection Schemes

Two pornography detection approaches, Yahoo Detector [6] and NuDetective [30] were evaluated using the built datasets. The evaluated techniques were executed on a desktop computer equipped with an Intel i7 CPU (quad-core), 16 GB of memory, and an Nvidia Titan-XP GPU.

Table 1 presents the throughput and detection accuracy results for Yahoo Detector and NuDetective. It can be seen that NuDetective significantly outperforms Yahoo Detector in terms of throughput, reaching up to 411 FPS, whereas Yahoo Detector reaches only 3.67 FPS. However, Yahoo Detector has a significantly improved detection accuracy, reaching 45.68% and 21.64% smaller FP rates for Pornography-2k and UCF101, respectively, when compared to NuDetective. In contrast, for the FN rate, NuDetective outperforms Yahoo Detector by 2.47% because NuDetective performs detection based on skin tone. Consequently, the detection throughput is significantly improved, but the FP rates are also considerably higher than that of a CNN-based approach.

Table	1: 4	Accuracy	and	through	out of	nornogra	nhv	detection	annroa	ches
1 and	1. 1	ACCUI acy	anu	un ougn	Jui	pornogra	pny	utitution	approa	ciics

	Detection Throughput	Dataset			
Detection Technique	(Frames per Second)	Pornography-2k		UCF101	
		FP (%)	FN (%)	FP (%)	
Yahoo Detector [6]	3.67	12.80	5.53	1.61	
NuDetective [30]	411.64	58.48	3.06	23.25	

### 3.3 Discussion

The evaluation of current pornographic detection techniques for videos shows that they cannot handle real-time detection even when performed using dedicated hardware. These techniques either have a low throughput or significantly high error rates. Running them on resource-constrained devices, environments in which pornographic content is often shared and watched, is not feasible owing to low accuracy or high computational demand. If they have a high FP rate, the user will no longer trust the detection mechanism—recall that the best FP rate was obtained using Yahoo Detector (Table 1). In addition, as these techniques classify an entire video as pornographic or normal, an FP causes the user not to be able to watch the entire video.

### 4 PRIVATE PARTS CENSOR

To address the real-time detection requirement and the lack of reliability in the classification of pornographic content in videos, we present the Private Parts Censor (PPCensor), whose objective is real-time detection and ease of use without the need for additional processing on the end-user's device. PPCensor runs on dedicated and general-purpose hardware, and the process of detection is transparent to the end-user. PPCensor's operation proceeds in two main stages: video-streaming query and handling and pornographic object detection.

The first stage addresses device query parsing, video download, video filtering, and related filtered responses. The process is entirely executed on a proxy server. Therefore, the only requirement is that the end-user configure the device settings to redirect network queries to the PPCensor server. The detection process does not require additional device processing and can be readily executed even on resource-constrained devices. For instance, a user responsible for an underaged user can configure the device that they want to protect.



Figure 1: PPCensor processing architecture.

The second stage is responsible for presenting a user-friendly obfuscation approach. PPCensor is the first approach that addresses the task of pornography detection as an object detection problem. It is based on the insight that a private part (explicit nudity) can be

identified as an object and properly filtered (obfuscated) inside a video frame. As a result, even if a misclassification occurs, the user experience is not significantly degraded because PPCensor deals with inappropriate content by obfuscating (blurring) private parts instead of blocking the entire video.

PPCensor processing architecture is illustrated in Figure 1. These two stages, including the architectural components, object detection techniques applied to the detection of pornographic video frames and the implementation of the proposed prototype, are detailed in the following subsections.

### 4.1 Pornographic Object Detection

Before examining the details of PPCensor's processing architecture, we describe its novel detection approach. Unlike related works that address the detection of pornographic videos through the analysis of entire video frames, PPCensor addresses the detection of pornographic content as an object detection problem. Therefore, PPCensor provides two essential benefits: improved detection throughput and easy-to-use detection. First, the detection of image objects can be achieved in near real time because several CNN architectures are used to handle object detection tasks as classification problems, significantly increasing the detection throughput. Second, filtering specific regions of the image (objects) instead of the entire image substantially improves the user experience because if an FP occurs, only part of the video image is obfuscated, instead of the entire video being blocked.

PPCensor considers each private part, either of a female or male, as an object for detection. Thus, the objective of PPCensor is to recognize four classes of specific objects considered Private Parts: male and female sexual organs (penises and vaginas), female breasts, and buttocks. It is important to note that these private parts are not shown in normal video streams, but they are often featured in pornographic videos. For example, in a single frame, an explicitly pornographic video can show all these private parts.

When considering private parts as objects for detection, PPCensor can detect pornographic videos using object detection techniques based on CNN. Videos can come from any URL on the internet.

### 4.2 Processing Architecture

The objective of using PPCensor is to perform object detection on pornographic content in real time without any additional processing on the user's device. Therefore, PPCensor is implemented as a proxy server, whose processing architecture is illustrated in Figure 1.

The architecture considers a set of user devices for which video media needs to be evaluated and filtered. For this purpose, the user must only configure their device to redirect the queries to the PPCensor proxy server using it as a web proxy. Therefore, all streaming video queries made by the device are sent to PPCensor. The PPCensor server, in turn, receives the queries through the request handler module, responsible for receiving and analyzing URL (Uniform Resource Locator) queries. As such, the purpose of the module is to establish a video query URL for further processing.

Subsequently, the content handler module queries the video streaming server for the related video frames. It downloads the streaming frames to the proxy server and forwards it to the next module, the content evaluator. This module, in turn, performs the proposed pornographic object detection and appropriately labels all identified image objects (described in Section 4.1). Finally, the content filter module filters the frames in real time by, for example, blurring the identified private parts within each frame.

## **4.3 PROTOTYPE**

A PPCensor prototype was implemented and deployed in a multithreaded process, as shown in Figure 2. The prototype considers three main entities: the end-user device, a proxy server (an implementation of PPCensor), and a video streaming server.



Figure 2: PPCensor prototype implementation architecture.

The video URL query process is responsible for the video query in PPCensor. In a production deployment configuration, a video URL query is represented by a request from the end user's device. To that end, a Python program was implemented to perform various video URL queries. All conducted queries are forwarded to the PPCensor (proxy) server implemented using Squid Proxy [46], a well-known network proxy server. The video URL lookup node is configured to use PPCensor as its proxy server.

The PPCensor server is implemented using three main processes. First, the proxy is implemented using Squid Proxy [46] version 4.5.

This tool is available to the public and is widely used as a web proxy on Linux-based servers. Over time, Squid Proxy logs all queries in a log file, which is analyzed and interpreted by another process, i.e., the query filtering and handling module. These processes are Python modules and continuously filter the Squid Proxy log. When a video URL is found, the prototype checks the URL to determine if it was already verified, using a cache-based blacklist and whitelist. If the queried URL was not verified, the corresponding media URL was downloaded from the video streaming server through the Pafy API [47], version 0.5.1. The process downloads a video at 0.5 FPS. A single frame is downloaded for each 2-second video interval, which is temporarily cached in a folder for further processing. Finally, the proposed pornographic object detection approach is implemented using Python in another process. This process runs the TensorFlow object detection API [48] using TensorFlow-GPU version 1.14. Thus, the module is executed continuously. When a new video frame is downloaded, the CNN model is applied, and the identified pornographic objects are labeled. The module also allows blacklisting of videos via the Squid Proxy blacklisting feature. Therefore, a pornography-related URL can also be blocked in real time, if required by parental control software, for example.

The video streaming server is responsible for answering both video URL queries and PPCensor queries. It is important to note that if the PPCensor server blacklist finds a pornography-related URL, the download from the blacklisted URL can be promptly interrupted by parental control. In this case, the end-user will no longer be able to access the video, even if the video stream is currently being downloaded. PPCensor, by default, blurs (obfuscates) only the private parts (objects) within a video frame.

## **5** EVALUATION

We present the PPCensor analysis by answering four research questions (RQ). (RQ1) Is it possible to treat the detection of a pornographic video as an object detection task? (RQ2) How challenging is the detection of each private part? (RQ3) How does PPCensor perform image classification tasks compared to state-of-the-art CNN? (RQ4) Does PPCensor facilitate real-time detection of pornographic video streams?

The following subsections outline the details of the dataset developed in our study. Additionally, the accuracy and throughput of PPCensor are evaluated.

### 5.1 Private Parts Object Dataset

In general, the data provided in the literature is labeled for the evaluation of schemes for detecting pornographic images or videos, rather than objects. Consequently, to evaluate the proposed PPCensor detection mechanism, we built a new dataset, the PPO (Private Parts Object Dataset).

The PPO dataset was based on the well-known Pornography-2k dataset [34]. Since we treat pornography detection as an object detection problem, we preprocessed the 1,000 pornographic videos in Pornography-2k. The preprocessing task filters the video frames according to their content. Thus, only explicitly pornographic images are selected. For this subset, we manually select the regions in each video frame that show a penis, a vagina, female breasts, or buttocks. To achieve this goal, we define a bounding box region that contains an object that is considered pornographic.

The final dataset consists of 50,870 pornographic frames, within which we selected 52,215 objects containing 13,607 penises, 11,346 vaginas, 13,963 female breasts, and 13,299 buttocks. Note that a single video frame may contain more than one target object. For the model building procedure, the dataset was divided into three parts: training, validation, and testing; the parts of the dataset were made using 60%, 20%, and 20% of the originally built dataset, respectively. Consequently, each part of the dataset is composed of mutually exclusive videos, each with their pornographic objects, thus allowing an adequate evaluation of the method. The training dataset is used for training purposes, the validation dataset is applied to the model training adjustment process, and the test dataset is used for the final evaluation of the model.

## 5.2 Accuracy of PPCensor

To evaluate the feasibility of PPCensor object-based detection, we first address RQ1, using the PPO dataset and several CNN-based object detection approaches. The evaluations were performed on a desktop computer equipped with an Intel i7 CPU (quad-core), 16 GB of memory, and an Nvidia Titan-XP GPU.

CNN Architecture	Detection Throughput (Frames per Second)	mAP (IoU 0.5)
Faster R-CNN Inception v2	10.87	63.50
Faster R-CNN NAS	1.19	56.53
Faster R-CNN ResNet 50	8.33	62.55
SSD MobileNet	59.55	55.92

### Table 2: Evaluation of PPCensor performance in pornographic object detection

All evaluated CNN architectures were trained using transfer learning [20, 21]. CNN used the weights obtained from the COCO [28] dataset, which is widely used for object detection purposes. Each CNN architecture was executed in 400,000 steps. Learning and decay rates were defined empirically for each architecture after several evaluations. The CNN-based techniques were evaluated according to

their detection throughput and mAP, as computed according to an IoU threshold of 0.5 to each object bounding box.

Significant differences between the accuracy (mAP) values were obtained for each evaluated CNN architecture, as shown in Table 2. It is important to note that the mAP considers an entire intersection of identified objects of at least 50% in this evaluation. However, in a production environment, the identified bounding box region can be expanded to improve the accuracy of the resulting model without a significant change in the user experience. The Faster R-CNN Inception's mAP architecture shows that it is possible to deal with streaming pornographic video as an object detection task.



Figure 3: Faster R-CNN Inception IoU and trade-off of inaccuracy.

In Figure 3, we investigate the relationship between the IoU threshold and the accuracy rates. It is possible to observe that varying the IoU threshold can improve the accuracy of object detection up to an IoU threshold of 0.8. However, increasing its value further does not significantly improve the true object ratio, while it does significantly decrease the object detection accuracy. Using a 0.8 IoU threshold value improves the object detection for up 3.84% when compared to a 0.5 IoU threshold.

To answer RQ2, we inspect the detection rate of each private part. We inspect the confusion matrix for the Faster R-CNN Inception architecture with an IoU of 0.8 (Operation Point, Figure 3), as shown in Table 3. It can be seen that private parts are generally not misclassified; about 11% of cases (on average) show as FNs. This means that CNN is able to identify which type of private part was found, although three objects have about 1.5% of false positives (on average): female breasts, vaginas, and penises.

Object Class		False Negative			
	Female Breasts	Penis	Vagina	<b>Buttocks</b>	(FN)
Female Breasts	2186	38	0	5	271
Penis	2	2236	11	4	247
Vagina	0	4	2199	39	258
Buttocks	5	7	38	2165	285

Table 3: Confusion matrix for Faster R-CNN Inception architecture (IoU 0.8).

Figure 4 shows examples of pornographic images labeled using the Faster R-CNN Inception architecture, in which it is possible to observe how the proposed object detection approach positively affects the user experience. The regions in the images with private parts are appropriately identified and obfuscated, and the user experience is not significantly degraded, because the user is still able to identify the context of the scene without being exposed to explicit nudity.

Figure 5 shows examples of non-pornographic images with misrecognized objects (FPs) when the Faster R-CNN Inception architecture is applied. Common examples of FPs include an arm or a finger misidentified as a penis (Figures 5-c, 5-d, and 5-h) and rounded regions being incorrectly labeled as female breasts (Figures 5-a, 5-e, and 5-g). In figure 5, it can be seen that the incorrect identification of an object does not significantly degrade the end-user experience because only a small part of the video frame is obfuscated. We pixelated people's faces in order to preserve their identification in figures 4-e, 5-a, 5-b, 5-c, and 5-d.

Under certain circumstances, the system parental control administrator may prefer to block the entire video, or parts of it, rather than obfuscating specific regions of the video frame (object) identified by PPCensor. For example, in the case of explicitly pornographic videos, it is preferable to block access to the entire video, instead of allowing access to filtered (obfuscated) video sequences. Therefore, to address RQ3, we also evaluated our proposed approach by applying it to image classification. PPCensor was customized to classify a video frame as pornographic if any private part was identified within the video frame. The PPCensor results were compared with the state-of-the-art CNN architectures used for image classification tasks.



Figure 4: Examples of private parts after identification by PPCensor with Faster R-CNN Inception (IoU 0.8).

The Pornography-2k dataset was divided into training, validation, and testing datasets, comprising 60%, 20%, and 20% of the original dataset, respectively. The CNNs were trained with transfer learning [20, 21]. In addition, a weight adjustment was applied using the validation dataset; finally, the final accuracy was measured using the testing dataset. The CNN architectures (CaffeNet [49], AlexNet [22], and Inception [19]) were executed for 400,000 steps. The learning and decay rates were empirically defined individually for each architecture after several evaluations. The obtained models were also evaluated using the UCF101 dataset; in this case, a pornography-classified video frame is considered an FP since the dataset consisted of only normal video frames.



Figure 5: Common FP examples from UCF101 dataset when using Faster R-CNN Inception (IoU 0.8).

Table 4 shows the PPCensor performance compared to traditional image detection techniques for the Pornography-2k and UCF101 datasets. Surprisingly, PPCensor outperforms all the evaluated state-of-the-art techniques for the Pornography-2k dataset in terms of FP, while also outperforming all other approaches except the Yahoo Detector for the UCF101 dataset.

PPCensor also shows similar results to other techniques regarding its FN rates, reaching 2.34% for the Pornography-2k dataset, an increase of only 0.31% compared with the best CNN architecture. However, the detection throughput of PPCensor is significantly worse than the other CNN architectures. It can be noticed that the detection performance of our proposed technique can be improved by applying other CNN architectures to detect objects (as shown in Table 2).

	Detection Throughput	Dataset			
Detection Technique	(Frames per Second)	Pornography-2k (Test Dataset)		UCF101	
	· · · ·	FP (%)	FN (%)	FP (%)	
Yahoo Detector [6]	3.67	12.80	5.53	1.61	
NuDetective [30]	411.64	58.48	3.06	23.25	
CaffeNet [49]	63.97	3.41	3.70	9.68	
AlexNet [22]	61.12	5.08	4.61	13.69	
Inception [19]	18.29	2.46	2.03	8.03	
PPCensor (Faster R-CNN Inception)	10.87	1.66	2.34	3.41	

Table 4: Comparison of PPCensor and CNN approaches when applied to image classification tasks.

The evaluation of video frames is only part of the PPCensor architecture (Figure 2). Therefore, despite the use of CNN with high detection throughput, the rest of the pipeline should also be improved. Specifically, real-time video download and filtering should be addressed. Therefore, in the next subsection, the overall throughput of the PPCensor architecture and the scalability of the connection for real-time video classification tasks are evaluated.

#### 5.3 PPCensor Performance and Scalability

We evaluated whether PPCensor can perform transparent and real-time detection of pornographic content using a proxy server, without any additional processing on the end-user device, to answer RQ4. The PPCensor node was equipped with an Nvidia Titan-XP GPU, which ran the Faster R-CNN Inception architecture with an mAP threshold of 0.8 (for implementation details, refer to Section 4.3).



Figure 6: PPCensor video download time.

Figure 7: PPCensor video classification time.

Figure 6 shows the trade-off between the video length and the video download time as a function of the number of connections. It can be observed that PPCensor can download a video for up to ten user connections in real time, as the download time is always shorter than the video length. In addition, the download time increases linearly for a single-user scenario. At the same time, there is a significant increase in execution time for ten users because the video download is a network- and CPU-bound. Since the evaluated node contains only eight threads, the execution time decreases significantly below a specific video length limit (20 min).

Figure 7 shows the trade-off between video length and classification time. It can be observed that the execution time increases almost linearly for both single-user and ten-user scenarios. This behavior is mainly because the classification task is GPU-bound. The evaluation of video frames is performed individually, rather than simultaneously, as only one model is loaded in memory in the prototype implementation.

Figure 8 shows the trade-off between video time and processing time for the entire PPCensor process, including the video download and frame classification. Detection can be performed in real-time for single- and ten-user scenarios because the process execution time remains shorter than the video length. In addition, the required processing time does not increase significantly compared to the classification time (Figure 7). This characteristic is due to the CPU-bound nature of the download process and the GPU-bound nature of the classification task. Therefore, the processes of download and classification do not hinder each other. As a result, PPCensor allows the detection of the pornographic video in near real time, without any additional processing on the end user's device.



Figure 8: PPCensor processing time.



Figure 9: PPCensor processing time for 5-min-long standardand high-quality videos. The proposed approach can handle, in real time, 35 users viewing standard-quality video sequences or 25 users viewing high-quality video sequences.

Finally, Figure 9 shows the relationship between the number of users and the process execution time for 5-minute streaming video scenarios. In these cases, PPCensor can accommodate, in near real time, up to 35 users watching a video in standard quality or up to 25 users watching a video in high quality. Thus, PPCensor transparently detects pornographic video streams in real time for up to 35 users without incurring any additional device processing demands. At the same time, the detection is performed in a user-friendly manner. It is important to note that the prototype implementation can be adjusted to improve observed issues, further increasing the detection throughput.

#### **6** CONCLUSION

Current publicly available tools for detecting pornographic content, when streaming videos do not enable transparent detection, are not easy to deploy in real-world environments. In this proposal, we present PPCensor, which is based on two main insights: detection and obfuscation of pornographic content. First, PPCensor treats the identification of pornographic content as an object detection problem. Such an approach provides user-friendly detection without significantly impacting the video watching experience even when an FP occurs or when scenes are showing private parts. An evaluation based on an analysis of more than 50 thousand video frames of (manually labeled) private parts revealed that the proposed technique could detect pornographic content in near real time. In addition, the technique generated similar results to those obtained from state-of-the-art proposals used for image classification tasks. Second, the PPCensor architecture is implemented as a video streaming proxy server; therefore, it does not incur additional processing on the end user's device, while performing transparent private part detection. An immediate positive impact of this approach is its usability on mobile devices such as smartphones and tablets. Our proposal is the first implementation that detects private parts as objects in real time for resourceconstrained end-user devices. The evaluation, which was carried out on a high-end desktop computer with a proxy server for video streaming, shows that PPCensor can accommodate up to 35 simultaneous connections from end-users in near real time. The PPO dataset and PPCensor source code are publicly available for download at https://seeplab.ppgia.pucpr.br/ppcensor.

#### ACKNOWLEDGMENTS

The authors thank the Brazilian National Council for Scientific and Technological Development (CNPq) for their partial financial support (grant nº 430972/2018-0), NVIDIA Co. for donating the Titan-XP GPU used in the experiments, and URCOP (*Unidade de Repressão aos Crimes de Ódio e à Pornografia Infantil*), a special division of the Federal Police of Brazil for pornography crackdowns, for their support to the project. We especially thank Marcelo da Silva Moreira, Rafaela Vieira Lins Parca, and Cassina Saad de Carvalho from the Federal Police for their extraordinary support of the project's objectives. Jackson Mallmann would like to thank the IFC (Federal Institute Catarinense) and Coordination for the Improvement of Higher Education Personnel (CAPES) for the scholarship, granting nº 231/2017.

#### REFERENCES

- [1] S. Statista, How Much of the Internet Consists of Porn? Available online: https://www.statista.com/chart/16959/share-of-the-internet-that-is-porn/. 2019. (Accessed 1 February 2020).
- [2] S. Pornhub, 2018 Year in Review. Available online: https://www.pornhub.com/insights/2018-year-in-review/. 2018. (Accessed 1 February 2020).
- [3] P.C.A. America, Prevent Child Abuse. Available online: https://preventchildabuse.org/. 2019. (Accessed 1 February 2020).
- [4] E. Enough, Pornography Statistics. Available online: http://enough.org/stats\_porn\_industry. 2019. (Accessed 1 February 2020).
- [5] M. Moustafa, Applying deep learning to classify pornographic images and videos. 7th Pacific-Rim Symposium on Image and Video Technology (PSIVT), 2015. Available online: https://www.researchgate.net/conference-event/PSIVT\_Image-and-Video-Technology\_2015/2645. (Accessed 1 February 2020).

- [6] J. Mahadeokar and G. Pesavento, Open sourcing a deep learning solution for detecting nsfw images. Available online: https://yahooeng.tumblr.com/post/151148689421/open-sourcing-a-deep-learning-solution-for. 2016. (Accessed 1 February 2020).
- [7] M. Perez, S. Avila, D. Moreira, D. Moraes, V. Testoni, E. Valle, S. Goldenstein, and A. Rocha, Video pornography detection through deep learning techniques and motion information. In Neurocomputing, Volume 230, Pages 279-293, 2017, ISSN 0925-2312. DOI: https://doi.org/10.1016/j.neucom.2016.12.017.
- [8] J. Wehrmann, G. S. Simões, R. C. Barros, and V. F. Cavalcante, Adult content detection in videos with convolutional and recurrent neural networks. Neurocomputing. Volume 272, Pages 432-438, 2018. DOI: https://doi.org/10.1016/j.neucom.2017.07.012.
- [9] X. Ou, H. Ling, H. Yu, P. Li, F. Zou, and S. Liu, Adult Image and Video Recognition by a Deep Multicontext Network and Fine-to-Coarse Strategy. ACM Trans. Intell. Syst. Technol. Volume 8, Issue 5, Article 68, 2017. DOI: https://doi.org/10.1145/3057733.
- [10] F. Idris and S. Panchanathan, Review of image and video indexing techniques. Journal of visual communication and image representation, Volume 8, Issue 2, Pages 146-166, 1997. DOI: https://doi.org/10.1006/jvci.1997.0355.
- [11] Y. Kuroki, T. Nishi, S. Kobayashi, H. Oyaizu, and S. Yoshimura, A psychophysical study of improvements in motion-image quality by using high frame rates. Journal of the Society for Information Display, Volume 15, 2007. DOI: https://doi.org/10.1889/1.2451560.
- [12] S. Wearsocial, The state of digital in April 2019: all the numbers you need to know. Available online: https://wearesocial.com/blog/2019/04/the-state-of-digital-inapril-2019-all-the-numbers-you-need-to-know. 2019. (Accessed 1 February 2020).
- [13] W. Rawat and Z. Wang, Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review. Neural Computing. Volume 29, Issue 9, Pages 2352-2449, 2017. DOI: https://doi.org/10.1162/neco\_a\_00990.
- [14] S. Pouyanfar, S. Sadiq, Y. Yan, H. Tian, Y. Tao, M. P. Reyes, M. Shyu, S. Chen, and S. S. Iyengar, A Survey on Deep Learning: Algorithms, Techniques, and Applications. ACM Computing Surveys (CSUR), Volume 51, Issue 5, Pages 1-36, January 2019. DOI: https://doi.org/10.1145/3234150.
- [15] M. Imani, D. Peroni, Y. Kim, A. Rahimi, and T. Rosing, Efficient neural network acceleration on GPGPU using content addressable memory. Proceedings of the 2017 Design, Automation and Test in Europe, Pages 1026–1031. DOI: https://doi.org/10.23919/DATE.2017.7927141.
- [16] H. Park, D. Kim, J. Ahn, and S. Yoo, Zero and data reuse-aware fast convolution for deep neural networks on GPU. 2016 International Conference on Hardware/Software Codesign and System Synthesis, CODES+ISSS 2016. (2016), Pages 1–10. DOI: https://doi.org/10.1145/2968456.2968476.
- [17] A. Karpathy, CS231n: Convolutional Neural Networks for Visual Recognition. Available online: http://cs231n.github.io/convolutional-networks/. 2016. Accessed: February 2020.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, 2016, Pages 770-778. 2016. DOI https://doi.org/10.1109/CVPR.2016.90.
- [19] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, Rethinking the inception architecture for computer vision. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Pages 2818–2826, 2016. DOI https://doi.ieeecomputersociety.org/10.1109/CVPR.2016.308.
- [20] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, Learning and transferring mid-level image representations using convolutional neural networks. In Proceedings IEEE conf. on computer vision and pattern recognition, 2014. DOI https://doi.org/10.1109/CVPR.2014.222.
- [21] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, A Survey on Deep Transfer Learning. 2018. DOI https://doi.org/10.1007/978-3-030-01424-7\_27.
- [22] A. Krizhevsky, I. Sutskever, and G. E. Hinton, ImageNet classification with deep convolutional neural networks. Communications of the ACM. 60, 6
- (May 2017), Pages 84–90. DOI: https://doi.org/10.1145/3065386.
- [23] Z. Zhao, P. Zheng, S. Xu, and X. Wu, Object Detection with Deep Learning: A Review. IEEE Transactions on Neural Networks and Learning Systems. 2019. Pages 1-21. DOI: https://doi.org/10.1109/TNNLS.2018.2876865.
- [24] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, and K. Murphy, Speed/accuracy trade-offs for modern convolutional object detectors. The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017. DOI https://doi.org/10.1109/CVPR.2017.351.
- [25] S. Ren, K. He, R. Girshick, and J. Sun, Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. IEEE Trans. Pattern Anal. Mach. Intell., Volume 39, Issue. 6, Pages 1137-1149, 2017. DOI https://doi.org/10.1109/TPAMI.2016.2577031.
- [26] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C. Fu, and A. C. Berg, SSD: Single Shot MultiBox Detector. Computer Vision -- ECCV 2016, Springer International Publishing. DOI https://doi.org/10.1007/978-3-319-46448-0\_2.
- [27] J. Redmon and A. Farhadi, Yolov3: An incremental improvement, CoRR, 2018.
- [28] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C.L. Zitnick, Microsoft COCO: Common objects in context. In ECCV, 2014. DOI https://doi.org/10.1007/978-3-319-10602-1\_48.
- [29] S. Karamizadeh and A. Arabsorkhi, Methods of pornography Detection: Review. In Proceedings of the 10th International Conference on Computer Modeling and Simulation (ICCMS 2018). ACM, New York, NY, USA, Pages 33-38. DOI: https://doi.org/10.1145/3177457.3177484.
- [30] M. Polastro and P. Eleuterio, Nudetective: a forensic tool to help combat child pornography through automatic nudity detection. In Proc. Workshop on Database and Expert Systems Applications (DEXA), 2010. DOI https://doi.org/10.1109/DEXA.2010.74.
- [31] R. Ap-apid, An Algorithm for Nudity Detection. Proceedings of the 5th Philippine Computing Science Congress, 2005.
- [32] F. Nian, T. Li, Y. Wang, M. Xu, and J. Wu, Pornographic image detection utilizing deep convolutional neural networks. Neurocomputing, 210 (2016), Pages 283-293. DOI: https://doi.org/10.1016/j.neucom.2015.09.135.
- [33] R. Gordon, A calculated look at fixed-point arithmetic. Embed. Syst. Progr., 11 (4), 1998, Pages 72-79.
- [34] D. Moreira, S. Avila, M. Perez, D. Moraes, V. Testoni, E. Valle, S. Goldenstein, and A. Rocha, Pornography classification: The hidden clues in video space-time, Forensic Science International. Volume 268, Pages 46-61, 2016. DOI: https://doi.org/10.1016/j.forsciint.2016.09.010.
- [35] S. Avila, N. Thome, M. Cord, E. Valle, and A.A. Araújo, Pooling in image representation: the visual codeword point of view. Computer Vision Image Understanding, Volume 117, Issue 5, 2013, Pages 453-465. DOI https://doi.org/10.1016/j.cviu.2012.09.007.
- [36] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, Going deeper with convolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Pages 1-9, 2015. DOI https://doi.org/10.1109/CVPR.2015.7298594.
- [37] P. Vitorino, M. Avila, S. Perez, and A. Rocha, Leveraging deep neural networks to fight child pornography in the age of social media. In Proceedings of Journal of Visual Communication and Image Representation, 2018. DOI https://doi.org/10.1016/j.jvcir.2017.12.005.
- [38] L. Wang, J. Liao, and C. Xu, Vehicle Detection Based on Drone Images with the Improved Faster R-CNN. In Proceedings of the 2019 11th International Conference on Machine Learning and Computing (ICMLC '19). ACM, New York, NY, USA, Pages 466-471. DOI: https://doi.org/10.1145/3318299.3318383.
- [39] G. Plastiras, C. Kyrkou, and T. Theocharides, Efficient ConvNet-based Object Detection for Unmanned Aerial Vehicles by Selective Tile Processing. In Proceedings of the 12th International Conference on Distributed Smart Cameras (ICDSC '18). ACM, New York, NY, USA. DOI: https://doi.org/10.1145/3243394.3243692.
- [40] B. Benjdira, T. Khursheed, A. Koubaa, A. Ammar, and K. Ouni, Car detection using unmanned aerial vehicles: Comparison between faster R-CNN and YOLOv3. In Proceeding of 1st International Conference Unmanned Vehicle Syst.-Oman (UVS), Feb. 2019, Pages 1–6. DOI https://doi.org/10.1109/UVS.2019.8658300.
- [41] P. Maheshwari, D. Alex, S. Banerjee, S. Behera, and S. Panda, Top View Person Detection and Counting for Low Compute Embedded Platforms. In Proceedings of the 2018 the 2nd International Conference on Video and Image Processing (ICVIP 2018). ACM, New York, NY, USA, Pages 35-43. DOI: https://doi.org/10.1145/3301506.3301548.
- [42] L. Zhang, Y. A. Zhang, Z. Zhang, J. Shen, and H. Wang, Real-Time Water Surface Object Detection Based on Improved Faster R-CNN. 2019. DOI: https://doi.org/10.3390/s19163523. https://www.mdpi.com/1424-8220/19/16/3523.
- [43] S. Saikia, E. Fidalgo, E. Alegre, and L. Fernández-Robles, Object Detection for Crime Scene Evidence Analysis Using Deep Learning. In: Battiato S., Gallo G., Schettini R., Stanco F. (eds) Image Analysis and Processing - ICIAP 2017. Lecture Notes in Computer Science, Volume 10485. Springer, Cham. DOI https://doi.org/10.1007/978-3-319-68548-9\_2.

- [44] R. Zhao, W. Ouyang, H. Li, and X. Wang, Saliency detection by multi-context deep learning. In Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR'15). 1265-1274. DOI: https://doi.org/10.1109/CVPR.2015.7298731.
- [45] K. Soomro, A. R. Zamir, and M. Shah, UCF101: A dataset of 101 human actions classes from videos. In the Wild. 2012. CoRR.
- [46] Squid Development Projects. Available online: http://devel.squid-cache.org/. (Accessed 1 February 2020).
- [47] Pafy. Available online: https://pythonhosted.org/pafy/. (Accessed 1 February 2020).
- [48] Tensorflow Object Detection API. Available online: https://github.com/tensorflow/models/tree/master/research/object\_detection/. (Accessed 1 February 2020).
- [49] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, Caffe: Convolutional Architecture for Fast Feature Embedding. In Proceedings of the 22nd ACM international conference on Multimedia (MM '14). ACM, New York, NY, USA, 675-678. DOI: https://doi.org/10.1145/2647868.2654889.
- [50] Y. Li, W. Ren, T. Zhu, Y. Ren, Y. Qin, and W. Jie, RIMS: A Real-time and Intelligent Monitoring System for live-broadcasting platforms. Future Generation Computer Systems. 87, Pages 259–266 (2018). DOI: https://doi.org/10.1016/j.future.2018.04.012.
- [51] L. Wang, J. Zhang, Q. Tian, C. Li, and L. Zhuo, Porn Streamer Recognition in Live Video Streaming via Attention-Gated Multimodal Deep Features. IEEE Trans. Circuits Syst. Video Technol. 1 (2019). DOI: https://doi.org/10.1109/TCSVT.2019.2958871.
- [52] S. Singh, A.B. Buduru, R. Kaushal, and P. Kumaraguru, KidsGUARD: Fine grained approach for child unsafe video representation and detection. Proc. ACM Symp. Appl. Comput. Part F147772, Pages 2104–2111 (2019). DOI: https://doi.org/10.1145/3297280.3297487.
- [53] M. Braun, S. Krebs, F. Flohr, and D.M. Gavrila, EuroCity Persons: A Novel Benchmark for Person Detection in Traffic Scenes. IEEE Trans. Pattern Anal. Mach. Intell. 41, Pages 1844–1861 (2019). DOI: https://doi.org/10.1109/TPAMI.2019.2897684.
- [54] L. Yang, Z. Qi, Z. Liu, H. Liu, M. Ling, L. Shi, and X. Liu, An embedded implementation of CNN-based hand detection and orientation estimation algorithm. Mach. Vis. Appl. 30, Pages 1071–1082 (2019). DOI: https://doi.org/10.1007/s00138-019-01038-4.
- [55] G. Guo, H. Wang, Y. Yan, J. Zheng, and B. Li, A fast face detection method via convolutional neural network. Neurocomputing. (2019). DOI: https://doi.org/10.1016/j.neucom.2018.02.110.



Jackson Mallmann



Eduardo Viegas



Altair Olivo Santin



**Roger Robson dos Santos** 



Jhonatan Geremias

- It is provided a publicly available dataset with four annotated private parts of a human body (objects);
- The dataset is the first publicly available dataset in the literature with annotated private parts;
- It is presented a novel pornography content detection technique that handles pornography detection as an object recognition problem.
- The proposed approach, the first in the literature, can perform the detection of private parts in real time with similar accuracy rates compared to traditional image-based state-of-the-art approaches;
- It is presented a novel CNN-based detection architecture (PPCensor) that is suited to the resource-constrained nature of the devices commonly used for viewing pornographic videos;
- The architecture implements the proposed object recognition process in a video-oriented streaming proxy server. Therefore, it does not require device modifications, either additional processing on the device, and executes remotely and transparently to the end-user;

Н

**Jackson Mallmann** received the BS degree in computer science in 2004 and the MSC degree in computer science in 2011 from PUCPR. He is currently working towards his PhD degree in computer science at PUCPR. His research interests include machine learning, pornography detection and computer security.

Altair Olivo Santin received the BS degree in Computer Engineering from the PUCPR in 1992, the MSc degree from UTFPR in 1996, and the PhD degree from UFSC in 2004. He is a full professor of Graduate Program in Computer Science (PPGIa) and head of Security & Privacy Lab (SecPLab) at PUCPR. He is a member of the IEEE, ACM, and the Brazilian Computer Society.

**Eduardo Viegas** received the BS degree in computer science in 2013, the MSC degree in computer science in 2016 from PUCPR, and the PhD degree from PUCPR in 2018. He is an associate professor of Graduate Program in Computer Science (PPGIa). His research interests include machine learning, network analytics and computer security.

**Roger Robson dos Santos** received the BS degree in information systems from PUCPR in 2014 and is currently working toward the MS degree at PUCPR. His research interests include machine learning and computer security.

**Jhonatan Geremias** received the BS degree in computer science from PUCPR in 2016 and is currently working toward the MS degree at PUCPR. His research interests include machine learning, pornography detection and computer security.

### AUTHORSHIP STATEMENT

Manuscript title: PPCensor: Architecture for Real-time Pornography Detection in Video Streaming

All persons who meet authorship criteria are listed as authors, and all authors certify that they have participated sufficiently in the work to take public responsibility for the content, including participation in the concept, design, analysis, writing, or revision of the manuscript. Furthermore, each author certifies that this material or similar material has not been and will not be submitted to or published in any other publication before its appearance in the *Future Generation Computer Systems*.

#### Authorship contributions

Please indicate the specific contributions made by each author (list the authors' initials followed by their surnames, e.g., Y.L. Cheung). The name of each author must appear at least once in each of the three categories below.

#### Category 1

Conception and design of study: J. Mallmann, A.O. Santin, E.K. Viegas;

acquisition of data: J. Mallmann, J. Geremias, R.R. dos Santos;

analysis and/or interpretation of data: J. Mallmann, E.K. Viegas, A.O. Santin.

### Category 2

Drafting the manuscript: J. Mallmann, E. K. Viegas;

revising the manuscript critically for important intellectual content: J. Mallmann, E.K. Viegas, A.O. Santin.

#### Category 3

Approval of the version of the manuscript to be published (the names of all authors must be listed):

Jackson Mallmann, Altair Olivo Santin, Eduardo Kugler Viegas, Jhonatan Geremias, Roger Robson dos Santos;

### Acknowledgements

All persons who have made substantial contributions to the work reported in the manuscript (e.g., technical help, writing and editing assistance, general support), but who do not meet the criteria for authorship, are named in the Acknowledgements and have given us their written permission to be named. If we have not included an Acknowledgements, then that indicates that we have not received substantial contributions from non-authors.

This statement is signed by all the authors (a photocopy of this form may be used if there are more than 10 authors):

Author's name (typed)	Author's signature	Date	
Jackson Mallmann	Judion Mol	lmann	March 6, 2020
Altair Olivo Santin	Donti		March 6, 2020
Eduardo Kugler Viegas	Where 3		March 6, 2020
Roger Robson dos Santos	Roy R. don So	nta	March 6, 2020
Jhonatan Geremias	Gerandan	premios_	March 6, 2020

# Conflicts of Interest Statement

Manuscript title: PPCensor: Architecture for Real-time Pornography Detection in Video Streaming

The authors whose names are listed immediately below certify that they have NO affiliations with or involvement in any organization or entity with any financial interest (such as honoraria; educational grants; participation in speakers' bureaus; membership, employment, consultancies, stock ownership, or other equity interest; and expert testimony or patent-licensing arrangements), or non-financial interest (such as personal or professional relationships, affiliations, knowledge or beliefs) in the subject matter or materials discussed in this manuscript.

#### Author names:

Jackson Mallmann Altair Olivo Santin Eduardo Kugler Viegas Roger Robson dos Santos Jhonatan Geremias

The authors whose names are listed immediately below report the following details of affiliation or involvement in an organization or entity with a financial or non-financial interest in the subject matter or materials discussed in this manuscript. Please specify the nature of the conflict on a separate sheet of paper if the space below is inadequate.

Author names:

This statement is signed by all the authors to indicate agreement that the above information is true and correct (a photocopy of this form may be used if there are more than 10 authors):

Author's name (typed)

Author's signature

Date

Jackson Mallmann

Altair Olivo Santin

March 6, 2020

March 6, 2020

March 6, 2020

Roger Robson dos Santos

Eduardo Kugler Viegas

Jhonatan Geremias

March 6, 2020

day Sonta

March 6, 2020