Designing Light Filters to Detect Skin Using a Low-powered Sensor

Muhammad Uzair Tariq
University of Central Florida

DESIGNING LIGHT FILTERS TO DETECT SKIN USING A LOW-POWERED SENSOR

by

MUHAMMAD UZAIR TARIQ
B.S. Lahore University of Management Sciences 2013

A thesis submitted in partial fulfilment of the requirements
for the degree of Master of Science
in the Department of Computer Science
in the College of Engineering and Computer Science
at the University of Central Florida
Orlando, Florida

Summer Term
2017

Major Professor: Pamela Wisniewski

ABSTRACT

Detection of nudity in photos and videos, especially prior to uploading to the internet, is vital to solving many problems related to adolescent sexting, the distribution of child pornography, and cyberbullying. The problem with using nudity detection algorithms as a means to combat these problems is that: 1) it implies that a digitized nude photo of a minor already exists (i.e., child pornography), and 2) there are real ethical and legal concerns around the distribution and processing of child pornography. Once a camera captures an image, that image is no longer secure. Therefore, we need to develop new privacy-preserving solutions that prevent the digital capture of nude imagery of minors. My research takes a first step in trying to accomplish this long-term goal: In this thesis, I examine the feasibility of using a low-powered sensor to detect skin dominance (defined as an image comprised of 50% or more of human skin tone) in a visual scene. By designing four custom light filters to enhance the digital information extracted from 300 scenes captured with the sensor (without digitizing high-fidelity visual features), I was able to accurately detect a skin dominant scene with 83.7% accuracy, 83% precision, and 85% recall. The long-term goal to be achieved in the future is to design a low-powered vision sensor that can be mounted on a digital camera lens on a teen's mobile device to detect and/or prevent the capture of nude imagery. Thus, I discuss the limitations of this work toward this larger goal, as well as future research directions.

This thesis is dedicated to my family. To my parents who provided me with good education and supported me in my hard times and stuck with me through everything, thick and thin. To my teachers who gave me all the knowledge and wisdom to achieve whatever I have achieved till now.

ACKNOWLEDGMENT

First and foremost, I would like to express my extreme gratitude to my advisor, Dr. Pamela Wisniewski, for her exceptional guidance and the encouragement she gave me throughout. This thesis would not have been completed without her constant support. She was always there to assist me and provide me with her excellent insight into the research work for thesis.

Secondly, I would like to thank Dr. Sanjeev Koppal to assist me with the development and experimental phase and providing me with the required equipment for the experiments.

Last but not least, I am grateful to the Socio-Technical Interaction Research (STIR) Lab team whose members offered me instant help whenever I needed with this project.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1: INTRODUCTION

The advent of mobile smart devices, digital image capture, and multi-media messaging services (MMS), has created a phenomenon known as "sexting [61]," which now places adolescents at risk of long-term repercussions, as such ephemeral exploration can now be immortalized forever in the digital realm. Prolific sharing combined with the permanence of digitally captured nudity is particularly problematic as dissemination of child pornography (e.g., naked imagery or sex acts involving a minor) is a crime punishable by law [46]. Even more concerning, however, is the prolonged effects of a momentary mistake, including sexual predation, emotional trauma, cyberbullying, and even suicidal behaviors that have been documented in these situations [62]. To prevent such acts of momentary mistakes, like taking a sexually explicit picture of someone (even with their consent) or taking one's own sexually explicit picture through a mobile camera, we must devise a scheme that senses these actions beforehand and gives teens a chance to rethink their decision. Detecting sexting behavior at this level would serve to combat the problem at the source, instead of once the damage has been done.

The goal of this work is to take a small, but necessary, first step toward a more cohesive solution to this larger problem. If we were able to detect risky online behaviors (e.g., a teen taking a nude photo or streaming video while unclothed) using the device teens use to connect to the internet (e.g., mobile smartphone, tablet, or laptop), then we would be able to mitigate these risks in more meaningful ways. Unfortunately, nudity detection in itself poses additional risks to teens, as a high-fidelity digitized nude image of a minor (possibly transmitted to a server for additional processing) already negates our goal of preserving the privacy of minors. Therefore, an integral part of this long-term goal of detecting nudity prior to digital capture is a sensor that integrates directly with a mobile application to decouple skin detection (performed by the sensor) from risk mitigation strategies (managed by the application layer) so that parents can customize how to handle problematic behavior based on the age and unique needs of their teen.

I designed four light filters and paired them with a vision sensor one-by-one to detect skin patterns in light spectrum that is incident on the sensor through the filters. First step was to learn the skin pattern and for that ten participants took 300 observations. Therefore, a dataset associated with real scenes was created, including 150 non-skin dominant setting (e.g., a landscape or picture of an object) and 150 skin-dominant settings (e.g., an almost nude "selfie" taken in a bathroom or a close-up of one's face). I tested four machine learning algorithms to determine the accuracy in which they could detect skin dominance based on the sensor output. Simple Tree offered an accuracy of 83.7%, precision value of 0.83 and recall value of 0.85. The long-term goal to be achieved in the future is that a low-powered vision sensor is part of a teen's mobile device that operates all the time and detects and/or prevents the digital capture of nude imagery through regular camera.

1.1 Problem Definition: Risks Associated with Digitized Nudity

A recent study claims that 15% of teens on Snapchat claim to have received sexually explicit photos and 4% of cell-owning teens ages 12-17 say they have sent sexually suggestive nude or nearly nude images of themselves to someone else via text messaging [47]. Teen sexting behaviors are perpetuated by mobile technologies and by several direct messaging applications that are available on smartphones like Kik, Snapchat, and AskFM [58]. Unfortunately, such activities often fall under the jurisdiction of child pornography laws. Child pornography is illegal, and the federal law has stated the following; "A picture of a naked child may constitute illegal child pornography if it is sufficiently sexually suggestive. Additionally, the age of consent for sexual activity in a given state is irrelevant; any depiction of a minor under 18 years of age engaging in sexually explicit conduct is illegal." [46]. So, while sexting behaviors may seem innocent and exploratory to teens, in reality, they may have severe negative consequences.

1.2 The Current State of the Art in Nudity and Skin Detection

Over the last decade, computer scientists have tried to address these problems indirectly through exploring various nudity and skin detection techniques. For instance, Deselaers [33] developed a method for detecting adult nudity in videos based on a bag-of-visual-features representation for frames; Kovac et al. [56] proposed a method for detecting skin color based on RGB color space. Other researchers have used image databases to train a classifier. Lin et al. [57] used a Support VectorMachine (SVM), which has learning skills in the image detection of human nudity. Amato [40] created an application, which is able to intercept images received through various communication channels (e.g., Bluetooth, MMS) on mobile devices based on the Symbian™ operating systems. Once intercepted, the images are analyzed by the component of the system that automatically classifies images with explicit sexual content. Commercially, Facebook, Twitter and Bing are working closely with organizations, such as the National Center for Missing and Exploited Children's *CyberTipLine Child Victim Identification Program*, to track down illicit photos of minors [49]. PhotoDNA [50], which was developed by Microsoft, seems to be the most popular and latest technological solution for detecting digital nudity by analyzing digital imagery and metadata compared to a database of known images.

The approaches taken across academic researcher and industry have both had their limitations. The commercial solutions primarily focus on preventing the dissemination of child pornography (which already exists), and therefore, the problem in a post-hoc fashion. I argue that a more effective approach is to curtail the problem at the source and take suitable mitigation approaches to prevent the creation and dissemination of such imagery in the first place. Otherwise, all of the computational academic work related to the detection of nudity has been done at the software-level only; a myriad of algorithms have been designed by the computer science community to increase the accuracy and efficiency of detecting nudity [4, 10, 14, 16, 21, 22, 29, 30, 39], but they all operate on the already digitized images, that is, digitally-stored instances of nudity. There is no published work articulating the need to reduce teen sexting behavior at hardware level.

Skin detection at the pre-digitization level ensures privacy and in the future in can be combined with detection of other spatial and/or temporal features in nude scenes to prevent teen *sexting*.

1.3 Research Overview

In Chapter 2, I present a comprehensive literature review of algorithmic, or otherwise computational, approaches to detect nudity and/or skin that could be used as a means of preventing teen sexting behaviors. In Chapter 3, I explain in detail the implementation of skin detector. First, I go into the detail of hardware including the specific design of filters. Then, I explain the setting of the experiment that I performed to test the accuracy of the skin detector. In Chapter 4, I present and interpret the results of the experiment. In Chapter 5, I discuss the implications of my work, the limitations associated with the implementation, and the future work that needs to be done in order to achieve the long-term goal of reducing teen sexting behavior at a more secure level. Finally, in section 5.3, I conclude my thesis acknowledging the work that has been done and the work needs to be done.

CHAPTER 2: BACKGROUND

To motivate my research, I performed a comprehensive review of the literature and identified 45 peer-reviewed articles that summarize the state-of-the-art in nudity detection approaches. Below, I describe how I performed the systematic literature search and synthesized the literature.

2.1 Systematic Literature Search

I searched for articles involving adolescent sexting behaviors from mobile smart phones (both through photos and video imagery), through various libraries and databases including; IEEE Xplore Digital Library, ACM Digital Library, and Springer-link to ensure a comprehensive coverage of the existing literature. I also used Google Scholar to conduct a wider search to make sure that I include all the multidisciplinary peer-reviewed article I can possibly find on the internet. For this review, I focused on nudity detection using the following search terms: “teen nudity”, “adolescent nudity”, “nudity detection”, “skin detection”, “explicit content”, and “censor nudity.” I used the following inclusion criteria to evaluate if an article was relevant to my review: 1) The study was a peer-reviewed published work, 2) The study was published between 2008 and 2016, and 3) The study must suggest a technique to detect nudity.

2.2 Synthesizing the Literature

I identified 45 relevant articles and coded them based on the various dimensions summarized in *Table1*. I present a summary of the key findings from this literature review, which serve to motivate my proposed solution discussed later in Chapter 3.

Table 1: Structure Codebook

Dimension	**Description (Codes)**
Capture Type	Detection of nudity before or after capturing/digitization of image (PRE, POST)
Media Type	The technique mentioned is applicable on images, videos, both, or none (IMAGE, VIDEO, GEN)
Nudity Class	The technique detects the nudity only when the subject is completely nude, private parts are naked, or it's only non-nude sexually suggestive (COMP, PP, NN_SUGG)
Nudity Type	Detection of nudity of teen (under 18), or general nudity (TEEN, GENERIC)
Platform Type	Provide solution as general software and at web browsers, or as mobile apps and services (GEN, MOBILE)
Detection Approach	The implementation makes use of Machine Learning techniques, Computer Vision or Natural Language Processing (ML, CV, NLP)
Dataset Type	Online Open Source: The dataset is publicly available online, OR manually generated private (ON_OS, SELF_GEN)
Dataset Accuracy	If a dataset is used for learning and testing of explicit content, the accuracy of the labelled classes is verified by some authority or it is not sure (VERI, N_SURE)
User Study	Does the article based on a survey where a certain population is studied (YES, NO)
Risk Mitigation Approach	Article suggests an implementation of blocking or the reporting of the detected content (BLOCK, REPORT)

2.2.1 Capture and Media Type

Out of the 45 articles reviewed, all of them focused on post-digital capture imagery (i.e., "POST"), as opposed to pre-digital capture (i.e., "PRE"). This means that a digitized nude, image had to exist in order for the nudity detection approaches to be effective. Already digitized and stored images and videos can be classified very accurately through sophisticated techniques like extraction of low and high-level features [25]. This requires large processing time and space. However, when these techniques are used in real-time, it is crucial to follow time deadlines and memory constraints [60]. This indicates that detecting nudity live or in real-time applications like Skype is challenging and reliance on low-level features is more appropriate in these cases. The post-capture digital media analyzed included, images (66 % of articles), videos (25 % of articles), text (11 % of articles), and mixed media (i.e., “GEN”). The articles that analyzed images mostly used visual learning techniques like feature extraction to classify nude or non-nude image.

Though quite a few articles [10, 16, 21, 22, 29, 39] focused on detection of explicit content in videos, few made use of the temporal correlations inside video data. Alierza et al. [14] utilized different novel features for obscene video content recognition including spatial, spatiotemporal, and motion based features - using 3D skin volume method. Rehanullah Khan et al. [44] used the Viola-Jones object detection framework that works on real-time data to detect skin. Mateus de Castro et al. [30] considered the contribution of the percentage of explicit frames while classifying the video as pornographic. Apart from these three papers, all who indicated their focus on video did not actually implement a technique that made use of distinctive video properties. Hence, their techniques could also be applied to single frames that is same as detecting nudity in images.

Text, in most cases, was analyzed on the metadata attached to the image or video under consideration. 11% of the articles that we studied fall under two sub-categories: 1) Either they proposed a pure NLP approach [4] where YouTube videos were classified using metadata, 2) Or they focused on combining already available techniques to come up with the nudity filtering system [5] where Wijesinghe et al. proposed improvement in Parental Control and Filtering System combining techniques like site restrictions, denial of access to web proxy servers, identification of images containing nudity and control over uploads of images of people. We coded these two sub-categories as ‘GEN’ (general) for media type.

2.2.2 Nudity Class

We categorized the types of nudity detected in three classes as per their level of explicitness. 40.1% of the articles proposed a method that would only detect nudity if the image was completely nude. These techniques relied more on **skin detection** percentage rather than region based feature extraction, or they used some hardcoded high-level feature assumption like the use of navel recognizing process in [37]. 7.1% of the articles presented a solution that would detect exposed private parts of one's body even if the body was mostly covered. They focused more on high and low-level feature extraction, e.g., in [8]. Four of the articles had such a complex way of extracting features that they could almost be classified as using contextual learning for nudity. Therefore, they could detect if a naked body image was sexual or non-sexual, like in the case of a breastfeeding mother. Hakan Sevimli et al. [20] made use of 4 descriptors (feature extaction methods) to classify images in 5 classes of nudity: normal images (class 1), swimming suit images (class 2), topless images (class 3), nude images (class 4) and sexual activity images (class 5).

2.2.3 Nudity Type

70.5% of the articles focused on general (primarily adult) nudity detection, as opposed to specifically detecting the nudity of a minor. Of the 29.5% of articles that focused on teen nudity detection, they all detected age and nudity separately. To determine age, very few made use of the image itself; rather they used additional information like text form meta data or information entered directly by the uploader of the image. Articles that propose solution to general nudity detection in terms of age, mostly make use of two very general steps: Skin detection and Pornography detection [26, 27]. For detecting teen nudity, the third step 'Age detection' is added. Mateus de Castro Polastro et al. [9] makes it clear that there are some intrinsic characteristics that need to be considered to distinguish child nudity from adult nudity. For example, mostly the child sexual content doesn't contain explicit sex scenes so the motion (in video) can be distinguished for that of an adult porn scene and for the audio part, sound can also be distinguished since it is almost absent in the case of child porn video.

2.2.4 Platform Type

Even though mobile social networking and dating apps are a major platform for sexting [59], only 5% of the articles focused on detecting nude content specifically via mobile devices. While the software-based detection techniques that most of the researchers presented were general enough to be used via mobile platforms, none were tested to see if mobile processor had adequate processing power for the detection task. Two of the articles present a third-party mobile-application to detect explicit content. Yong Lin et al. [23] talks about an app that detects skin and classifies it, and for further detection of nudity the app would conduct a user poll. Giuseppe Amato et al. [40] presents an app that is more of a background service, an interceptor that detects nudity or other explicit content being received through MMS and Bluetooth messaging. Apart from these two articles, all others provide a general solution that can be implemented at different level of abstraction of the operating systems, depending on the authority. For instance, solutions could be implemented by the Internet Service Provider at the network level to provide information to law enforcement agencies or by an application that is controlled by parents of the user device.

2.2.5 Detection Type

From a computational perspective, most articles (86.4%) employed methods from the field of Machine Learning and Computer Vision (CV). Only 13% of the articles proposed Natural Language Processing techniques. For instance, Raverkar used NLP approach to classify YouTube videos [4]. A few papers presented pure CV solution without any learning aspect, for example, Pedro Ivan [6] used the RSOR algorithm that performed recognition and selection of the largest region in a segmented image. Some articles (e.g., Mateus de Castro Polastro et al. [7]) used mixed techniques from ML, CV, and NLP to come up with a comprehensive technique to detect nudity. Vanhove et al. [10] presented a solution that used Picture Analysis, Text Analysis and Audio/Video Analysis for social network monitoring.

2.2.6 User Study and Risk Mitigation

A common theme among the nudity detection approaches was that they came from a purely computational perspective and failed to incorporate any aspects of user-centered design or needs analysis. None of the articles included formative or summative user evaluations of the solutions developed. For instance, [52] talks about a parental control software and briefly mentions mitigation approaches that require a user study, which can consider different mitigation approaches to the problem depending on the relation or trust between parents and children. In terms of risk mitigation strategies for once nudity detection has occurred, the article set was also lacking. About 22.7% of the articles suggested naive approaches, such as blocking the explicit content or reporting it. In another 27% of the articles they suggested that their solutions could be used by law enforcement agencies, security companies, or parents. None of the articles suggested any kind of design that would directly engage teen users in a way to address the root of the risky behavior. In summary, the majority of the articles focused on risk detection over risk mitigation.

2.3 Gaps in the Literature

My review indicated some potential gaps in the literature that can inform new research directions, some of which I take in my own research design. Most articles made use of machine learning, and learning the features from images and videos and classifying them accurately requires more memory and processing time. This is the main reason that the articles that give a general solution to nudity detection problem do not specify the use of their method on mobile platforms because of the limited memory and processing power of the mobile device. Giuseppe Amato et al. [40] mentions that “Classifiers can recognize and discriminate between harmless and offensive multimedia contents (in a mobile device). However, the complexity of such systems discourages from implementing and running them on small mobile devices”. On the other hand, every device has the connectivity capabilities necessary for sending and receiving relatively rich amount of information which gives software developers the liberty to do extensive calculations at server side. However, interruption in connectivity could be a problem especially when dealing with real time detection.

Most of the articles that we studied propose skin and pornographic detection methods irrespective of the subject’s age. Articles that deal with teen nudity usually make use of textual data attached with the image to know the age of the subject. Few articles made use of the image itself to come up with an age through visual learning. For instance, almost all the detection techniques mentioned above either detect complete nudity or the private areas of a human body. Few papers [20, 25, 30] proposed methods that used contextual learning to detect image that are sexually suggestive. As far as adolescent online risk is concerned, there has not been a study to understand what level of nudity actually constitutes risk for teens. Therefore, detecting nudity in binary terms may not adequately serve to mitigate risks. Thus, more studies should focus on understanding the thresholds used to determine ground truth in nudity detection. As I discussed earlier, there was no article, to my knowledge, that suggested a design that would directly engage teen users in a way to address the root of the risky behavior. In short, the majority of the articles focused on risk detection over risk mitigation and failed to incorporate any aspects of user-centered design or a formative/summative user evaluations of the solutions developed. For capture and media type, the literature thus far has presented techniques to detect nudity in digitized images or frames of videos. However, a mechanism to detect nudity at pre-digitization level can be more useful in two major ways:

1) **Security**: To my knowledge, there is no peer-reviewed literature that deals with the issue of handling teen sexting behaviors before they have already been digitally captured via a nude image or video. At this point, the damage has already potentially been done. The New York Times reported that once an Apple iPhone, iPod Touch, or iPad owner grants permission for an application to access location information from their device, the application can potentially copy their photo library [51]. In several high-profile examples, celebrity photos have been leaked onto the Internet after their phones were hacked.

2) **Real-time detection**: At the heart of almost all techniques provided in the literature comes a basic step of feature extraction. So far, there has been more emphasis on *accuracy* of classification of nude images which means that techniques that have been proposed are extracting both low and high-level features. This requires considerable processing time and space, whereas, dealing with only low features will require less

processing time and that would help detecting nudity in real time with a compromise on accuracy. However, this way nudity detection in live or real-time video applications like Skype will be made possible.

2.4 Skin detection as an integral part to Nudity Detection

Detection of nude scenes requires detection of all kind of contextual and visual features in an image. In reviewing the literature, one of the important features for nudity detection was **skin** [64, 65]. Islam et al. in [2] states that, "Nudity and pornography have a direct link with human skin. In fact, no pornography can exist without exposure of human skin. Apart from pornography, a wide range of image processing applications exist, where skin detection is playing a crucial role. Using color as a detection cue has long being recognized as a robust feature and has become a popular choice in human skin detection techniques. Human skin has a characteristic color which is easily distinguishable from the colors of other objects."

Eight of the papers [2, 3, 18, 20, 27, 32, 44, 45] reviewed implemented skin detection as a part of the procedure to detect nudity. Islam [2] used Wavelet transform that involves recursive filtering and sub-sampling. It has discriminating ability in texture analysis, that facilitates capturing subtle differences between child and adult skin texture. K.K. Bhoyar [3] proposed three-layer feed forward neural network used for skin color classification with three neurons in input layer, five neurons in hidden layer and two neurons in the output layer. The two neuron in the output layer represents skin class and non-skin class. Digambar Povar et. al [18] and Wayne Kelly [32] used clustering in color space(s) to filter skin tone. Hakan Sevimli et al. [20] used an already implemented method based on inferring pixels on statistical skin and non-skin models which are represented and trained with Gaussian Mixture Models. Fernando Roberti de Siqueira [27] constructed color histograms through both the skin and non-skin groups of RGB images. They applied certain threshold on the histograms to classify pixels as skin and non-skin. Rehanullah Khan [44] used adaptive skin color modeling where pixels that are most likely non-skin are discarded from a detected region of pixels and the region is extracted for further processing. Bei-bei Liu [45] proposes to detect pornographic images in a two-stage scheme. The first step employs a content-based image retrieval technique (CBIR) to determine whether the image has human in it. The second step is a skin color model established to analyze the skin-like pixels and identify the presence of pornographic content.

Comparing the current research to that of above, I am dealing with pre-digitized image and hence I have an added advantage of processing continuous domain of light spectrum falling onto the sensor. We know that a digitized image through sampling loses some information. I, therefore, studied the pattern of electromagnetic spectrum reflected from a skin. Elli Angelopoulou et al. [63] performed an experiment to show that skin reflectance exhibits a “W” pattern in wavelength domain. I based my work on this paper by Elli and took their work further to design light filters to detect this pattern through a low-powered vision sensor.

2.5 The Big Picture versus Current Work

The motivation behind this research is combating the negative outcomes associated with adolescent “sexting” by addressing the problem at the source; the long-term goal is to use a low-powered light sensor that can operate 24/7 and can be mounted on a digital camera lens to detect and/or prevent the capture of nude imagery. I am using a low-powered sensor with four light-filters that are designed to detect skin. As already discussed in the previous section, skin detection is an integral part of detecting nude imagery. My goal is to detect risky sexting behaviors of teens and deter it, so that the risk associated with digitally captured nude imagery of minors is nullified. Therefore, there is an inherent trade-off between using a high-fidelity image to accurately extract visual features (thereby enhancing the risk I am trying to mitigate) and preserving teens’ privacy while finding a way that can still detect and mitigate risks. My rationale is that skin detection might be a reasonable compromise to this trade-off. In the upcoming chapter, I will discuss in detail how to design of the filters to be able to combine with the sensor to come up with a skin detector that detects skin at pre-digitization level.

CHAPTER 3: RESEARCH METHODS

3.1 Scope of this Thesis

This thesis implements a low-powered vision sensor prototype, combined with custom filters that I designed, to detect skin at the pre-digitization level with good accuracy. I created a dataset associated with real scenes, including 150 non-skin dominant setting (e.g., a landscape or picture of an object) and 150 skin-dominant settings (e.g., an almost nude "selfie" taken in a bathroom or a close-up of one's face). The output of the sensor was in the form of four real values (rounded down to integers), one for each filter. Then, I ran four machine learning algorithms i.e., Simple Tree, Logistic Regression, Linear SVM and Quadratic SVM, to determine the accuracy of skin detection. The prototype I designed (described below) will act as a valuable milestone towards the greater goals of this project.

3.2 Hardware Configuration

This section describes the physical hardware components and configuration of the prototype. Since the basic vision sensors (e.g., *FireFly* Vision Sensor) cannot be connected to a cell phone, they cannot be operated directly by a cell phone's processor to receive data. Therefore, for the sake of experiment I will use a separate processing chip i.e., *Arduino*, for testing purposes (connecting a processing chip to the vision sensor/chip to collect data and perform evaluation for the proposed solution).

3.2.1 Hardware Components

Processor: I am working with Arduino/Genuino Board (*Figure 14* in *Appendix*) using the open-source Arduino Software (IDE) which makes is easy to write code and upload it to the board. The environment is written in Java and based on Processing and other open-source software.

Vision Chip (sensor): I have FireFlyBig (See *Figure 1* and *Figure 12* in *Appendix*). The Firefly series of vision chips are a series of flexible resolution vision chips designed for a broad set of visual sensing applications. Two chips are available in this series, the FireflyBig having a resolution of 480x256 pixels, and the FireflySmall with a resolution of 128x256 pixels. I used the former one.

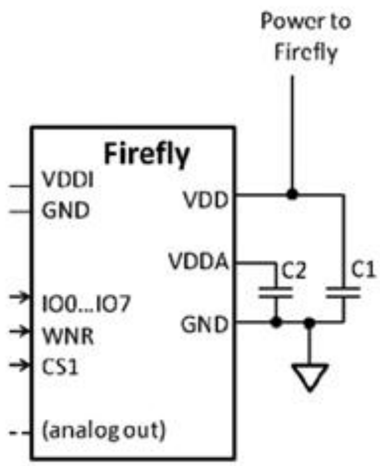

Figure 1: FireFly Chip Inputs and Outputs.
WnR: Write/ Not-Read PIO12B determines
direction od signals IO0…IO7, internally pulled down (read)

3.2.2 Hardware Configuration

The sensor is connected to a processor to send and receive signal. For our purpose, we don't need the 8-bit port. The analog output of the FireFly chip is providing us with pixel values of the pixel array, serially. The following (*Figure 2*) is a general configuration to connect them both.

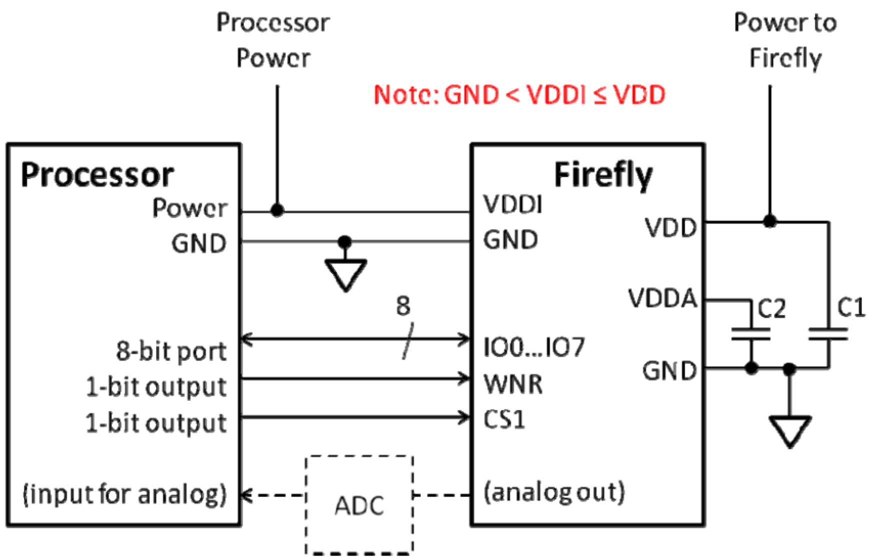

Figure 2:Processor connected to FireFly Sensor

<u>Process for configuring the hardware:</u>

- Connect the Arduino Board to the sensor and make sure that all the power and signal ports are connected correctly.
- Write a testing code (*Blink program*) on Arduino IDE, save it to the Arduino Board and run the code to see if everything is working properly. If the output LED on the Arduino blinks, that means that the Arduino is working properly.
- Run the code to get value from the sensor. Sensor can be tested by covering the pixel array with palm and then taking it away and see if the value that you get are considerably different. Use filters from SwatchBook (*Figure 13* in *Appendix*) to put between the sensor and the subject. Check to see if the value received changes with different filters.

These steps ensure that everything is connected and working properly and the equipment is ready for the main experiment. The following figure (*Figure 3*) shows an actual setup of the prototype where a green filter has been placed in front of the sensor.

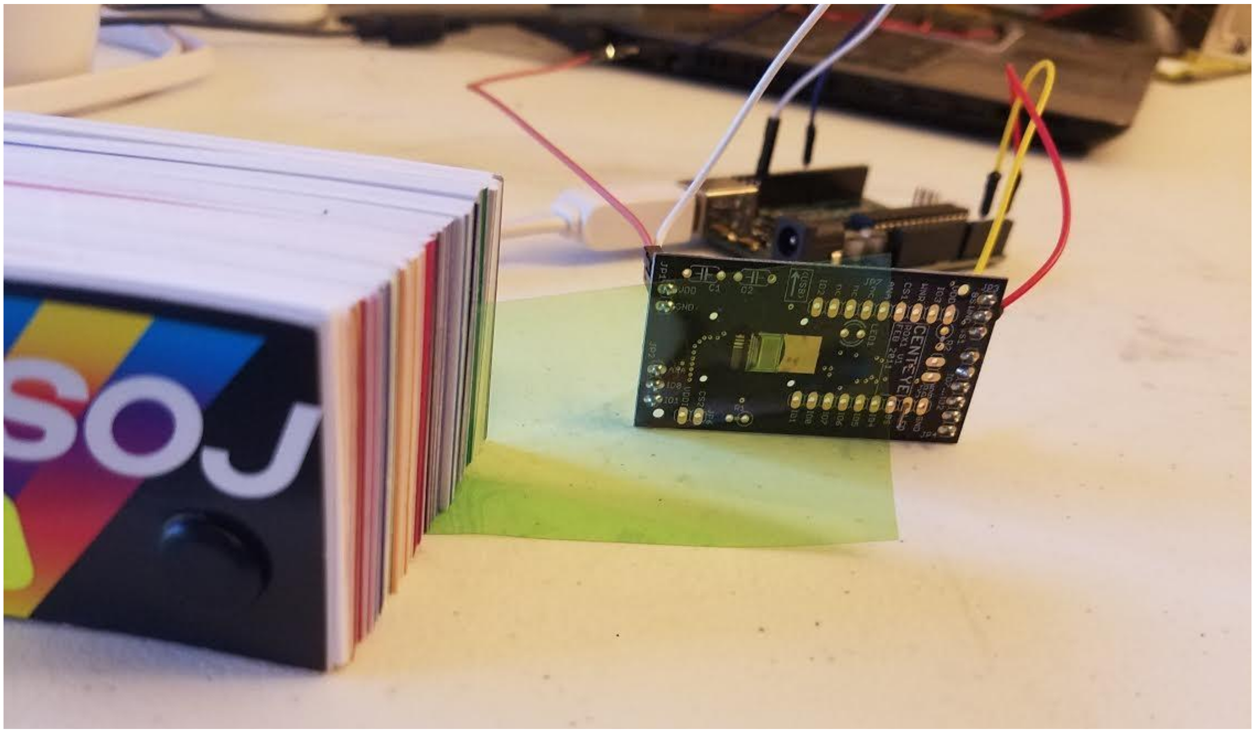

Figure 3: Setup of the Prototype

3.3 Designing Customized Light Filters

Before running an experiment, we need to understand what data we get from the sensor and how filters work on the sensor to provide the useful data that we need. I designed some specific filters to meet my requirements, using basic filters available in market.

3.3.1 Sensors, Sensor Data, and Light Filters

In my hardware configuration, I intentionally used a sensor with **no lens**. Therefore, the data is totally defocused. The sensor could act like a regular camera if a lens had been used to focus the light coming from the scene/object onto the pixel array. However, my aim was to defocus the image in order to maintain privacy. Defocusing the image means less feature detection, which is a key trade-off (between preserving privacy and accurate nudity detection) explicitly made by this research design. This also meant that that I could only detect skin, not actual nudity, using the sensor without a lens, which is further discussed in the limitations of this research. An ideal case for detecting nudity with privacy would be when the data that is captured on the pixel array is defocused enough not to be converted back to a recognizable nude image and at the same time has enough distinct features to detect nudity. The pixel array from the sensor provides almost the same value on every pixel for a single capture, which is then averaged to record one value per capture. The value comes out to be a real number greater than zero (represents light intensity in general). In order to allow different features to be present in the data (more than a single value), I took four different captures of the same object/scene with four different light-filters each. I designed these four filters.

3.3.2 Filter Design

The selection of the four filters (*Filters A, B, C, D*) that I designed can be justified by Elli Angelopoulou et al. [63]. According to the paper, the reflectances of various tones of human skin in the visible range of electromagnetic spectrum form a certain pattern. All the measured reflectances, except for the more dark-skinned (black) people, exhibit a localized "W" pattern (two dips with a bump in the middle) in the middle of the visible spectrum. Following graph (*Figure 4*) is taken from [63]. It shows the pattern of reflectances of various tones of human skin for wavelengths of visible light. Table 2 mentions the approximate position of the two peculiar dips and a hump in the graph.

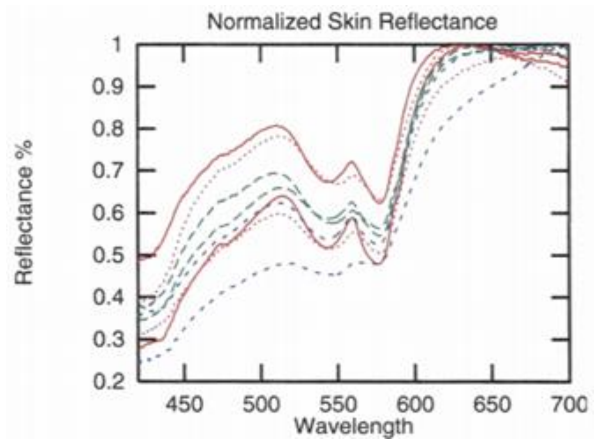

Figure 4: Reflectances of different tones of human skin in the visible range of EM spectrum

Table 2: Wavelengths of Local Min and Max

Feature	Location Median λ	Location Mean λ	Standard Deviation
Left Min.	546.42	546.56	2.54
Local Max	559.48	559.72	0.95
Right Min.	576.26	575.45	2.34

Ideally, the filters that I select should act as band pass filters allowing four roughly-distinct range of wavelengths to pass through to the sensor. The *Filter A* should transmit all the wavelengths from 500 to 550nm. The *Filter B* should be designed to transmit only the wavelengths from 550 to 565nm. The *Filter C* should pass wavelengths ranging from 565nm to 580nm and the *Filter D* should allows all the wavelengths above 580 to pass through. Ideal filter will pass the required range of wavelength with 100 percent transmission rate and stop the rest of the wavelengths completely. The following figure (*Figure 5*) shows the ideal versions of the four filters that we need to perform our experiment.

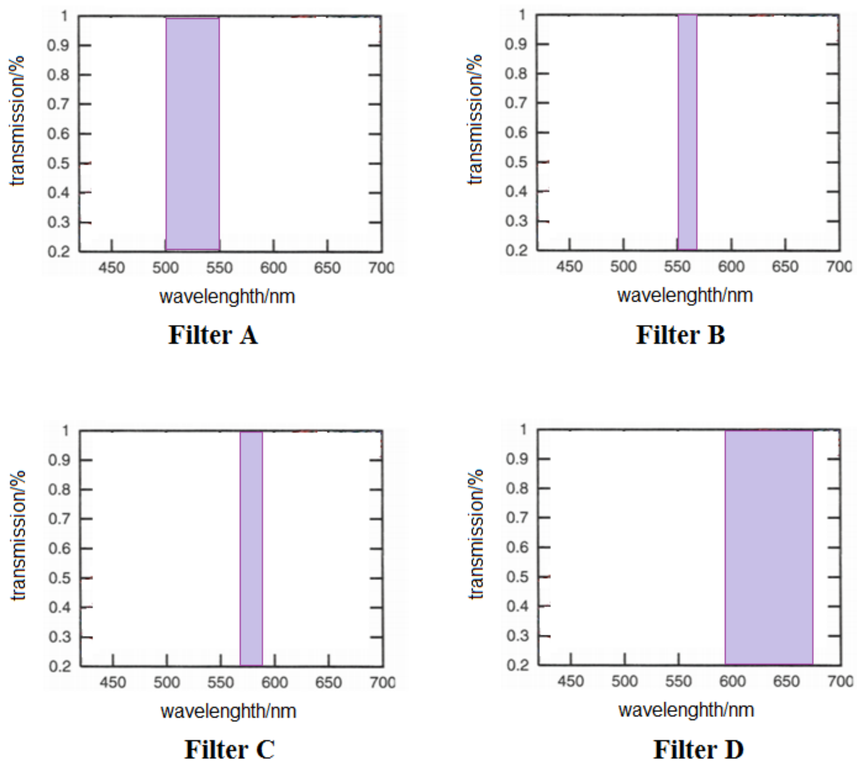

Figure 5: Ideal Light Filters needed for the Experiment

However, this is not possible practically. The band-pass filters that I designed, therefore, are not ideal. They roughly resemble the ideal filter. The filters available in the market do not bear any close resemblance to the ideal filters shown above especially the band-pass filters B and C. Therefore, I combined multiple filters to come up with the designs that resemble the above-shown filters. Simple optics and mathematics show that multiplication of spectral energy distributions (SEDs) of overlapping filters give the resultant SED of the combination of the filters. To illustrate this with an example, I will show how the SED curve of some commonly available filters can be used to design a required filter. The following figure (*Figure 6*) shows how the filters Deep Straw (Roscolux #15) and Leaf Green (Roscolux #386) can be combined to make the Filter B (in our case).

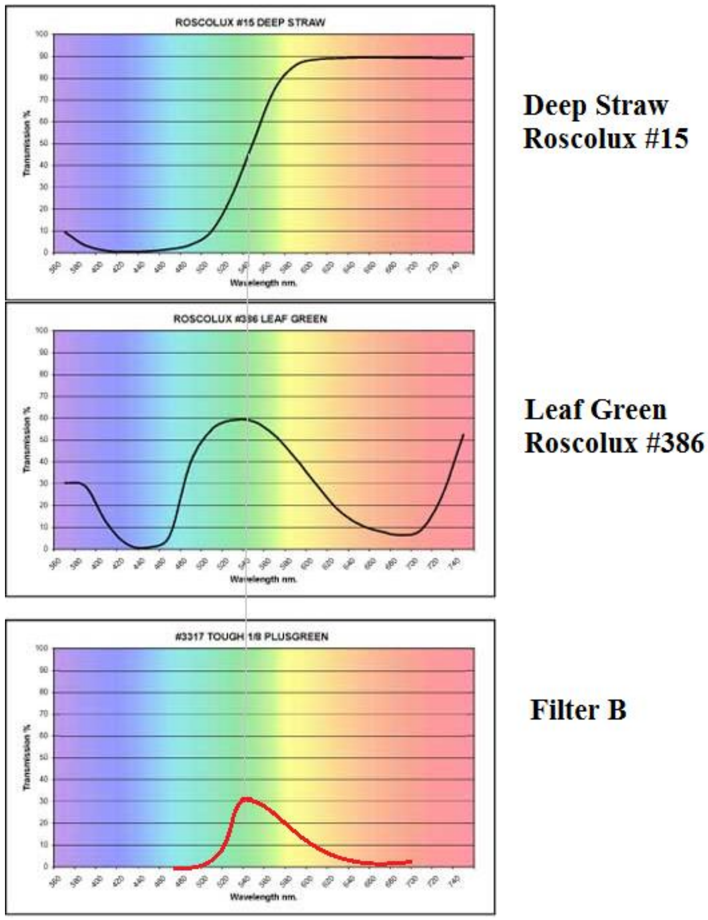

Figure 6: Mixing Deep Straw and Leaf Green to make Filter B

Similarly, for Filter 'A' I used Hemsley Blue and Moss Green, for Filter 'C' I used Apricot and Leaf Green, and for Filter 'D' I just used single filter i.e., Deep Amber. (Refer to *Figure 15* in *Appendix* for the SEDs of common filter mentioned here). These four filters differ from the ideal filters in bandwidth and percentage of transmission. They have low transmission and loose bandwidth partially allowing some of the undesirable wavelengths (compare Filter B in *Figure 5* and *Figure 6*). They can be improved given more variety of filters available in market.

3.4 Experimental Research Design

3.4.1 Study Overview

I used the FireFly Vision chip to capture 300 real scenes. Half of the scenes were category "1" which means that they were skin dominant while the other category (i.e., "0") contained regular scenes from indoors and outdoors. Skin dominant scenes were defined to be the ones where more than 50% of the scene in front of the sensor had skin (see *Figure 7* for an example of skin dominant and non-skin dominant scene). Since there was no lens, every pixel on the lens got the same value (light coming from everywhere) and hence that value was recorded. Each scene was captured four times with 4 different filters that we mentioned in the previous section. The set of these four values is actually one scene or one observation. The field of view of the sensor was restricted by covering the periphery of the pixel array with a hollow dull-black cylinder of radius 5 *mm* and a height of 20 *mm*. For training and testing purposes, the filters were switched manually. However, more detailed discussion on integration of the four filters can be found in *Section 5.1*.

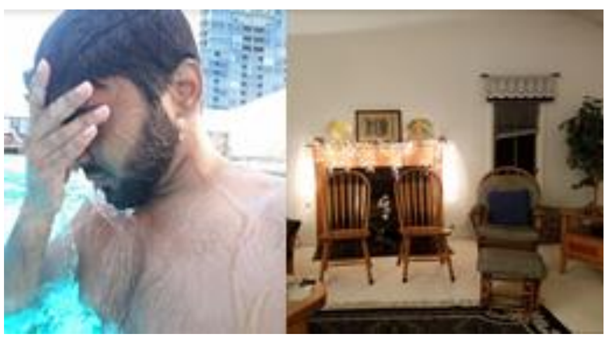

Figure 7: Left image is an example of skin dominant scene
Right image is an example of non-skin dominant scene

3.4.2 Creating a Training and Test Data Set

Ten individuals, including myself, volunteered to help create the data set. They used the vision sensor to capture 30 real scenes each (15 usual and 15 skin-dominant). Each scene was captured 4 times with the 4 different filters. Therefore, we created a total of 300 observations (10 people x 30 scenes = 300). Fifteen of the 30 scenes were taken with high skin content, exposing skin (non-nude, *see Section 5.1* to know why nude scene is not important) in their bedrooms and washrooms with subjective lighting (mostly fluorescent and incandescent, monochromatic lights were not used) and background. They usually took their selfies depicting typical 'nude selfies' taken by teens indoors. The participants were asked to take 15 skin dominant captures first and then 15 non-skin dominant captures. They subjectively chose 15 different scenes of **more than 50%** skin, recorded the readings with four filters and kept them aside. Then they chose different scenes of **less than 50%** skin and recorded the readings with four filters. Distance of sensor from the image subject does not matter but only the fraction of skin reflectance area in the field of view of the sensor. Participants were chosen based on their skin tones. They varied from brown, through beige, to white. I did not include a very dark-skinned person (black), since the "W" pattern for the reflectance of spectrum mentioned by Elli Angelopoulou et al. [63] does not exist for them.

3.4.3 Data Analysis Approach

The four values that we get using the filters make our feature vectors while the response variable is categorical (0 or 1). A total of 300 observations were split into test and training set, and were used to test the accuracy of four different classification algorithm. These algorithms (*Simple Tree, Logistic Regression*, *Linear SVM* and *Quadratic SVM*) are common for supervised learning tasks, and are chosen considering the type of data. The data that I have is simplified since filters at the hardware level already took care of much of the complexity of the task. The data is only 4 real-valued feature vectors that needs to undergo binary classification. I chose the above-mentioned algorithms based on a famous book by Alex Smola et al. [64]. Since I am detecting skin through only four filters, low accuracy is already expected (see *Section 5.2* on how to increase accuracy) and therefore achieving a classification accuracy of 80% should be considered a success. Skin detectors proposed in [2,3,18,20,27,32,44,45] got accuracy around 70% to 90%.

To perform the analysis, I used *Statistics and Machine Learning Toolbox* of MATLAB. The *Classification Learner* app is included in the toolbox that allows various training models to classify data using supervised machine learning. The toolbox also let us run these selected models in parallel and display the result in different tabs to compare. I tested four models based on the algorithms selected above and train the models with 250 observations (125 being skin-dominant). Then I test the models on 50 observations (25 being skin-dominant) to record the accuracy, precision and recall of the models. Results of the four models are provided in Chapter 4.

CHAPTER 4: RESULTS

4.1 Summary of Results

The results of the four models are shown in *Table 3* with figures that illustrate these of these models shown in *Figures 8-11*. To do the analysis, I made use of *Statistics and Machine Learning Toolbox* of MATLAB. The Toolbox has an app, *Classification Learner*, that offers numerous training models to classify data using *supervised* machine learning. Considering the simplicity of the data, I applied the following four classification models from the options: *Simple Tree*, *Logistic Regression*, *Linear SVM* and *Quadratic SVM*. I applied 6-fold cross validation to all model. The app takes care of the parameters itself. They are autotuned with optimal regularization constants and hence the app outputs the best accuracy possible for each model. I did not perform any analyses on individual skin tones as the numbers of sub-samples are too less to determine individual accuracies for different skin tones.

Table 3: Model Results

Model	**Accuracy**	**Precision**	**Recall**	**Prediction Speed**	**Training Time**	**Model Parameters**
Simple Tree	83.7%	0.83	0.85	~17000 obs/sec	0.96223 sec	Preset: Simple tree Maximum number if splits: 4 Split criterion: Gini's diversity index Cross Validation: 6-folds
Logistic Regression	77.3%	0.77	0.77	~6200 obs/sec	13.459 sec	Preset: Logistic Regression Cross Validation: 6-folds
Linear SVM Model	74.7%	0.77	0.71	~6600 obs/sec	3.36 sec	Preset: Linear SVM Kernel function: Linear Cross Validation: 6-folds Multiclass method: One vs One
Quadratic SVM Model	94.3%	0.94	0.95	~19000 obs/sec	1.1068 sec	Preset: Quadratic SVM Kernel function: Quadratic Cross Validation: 6-folds Multiclass method: One vs One

The following four plots exhaust all cases of true positives, true negatives, false positives and false negatives for each model. The accuracy of the last model is 94.3% and hence we can see in the plot that it has very less ‘dotted lines’ (incorrect cases).

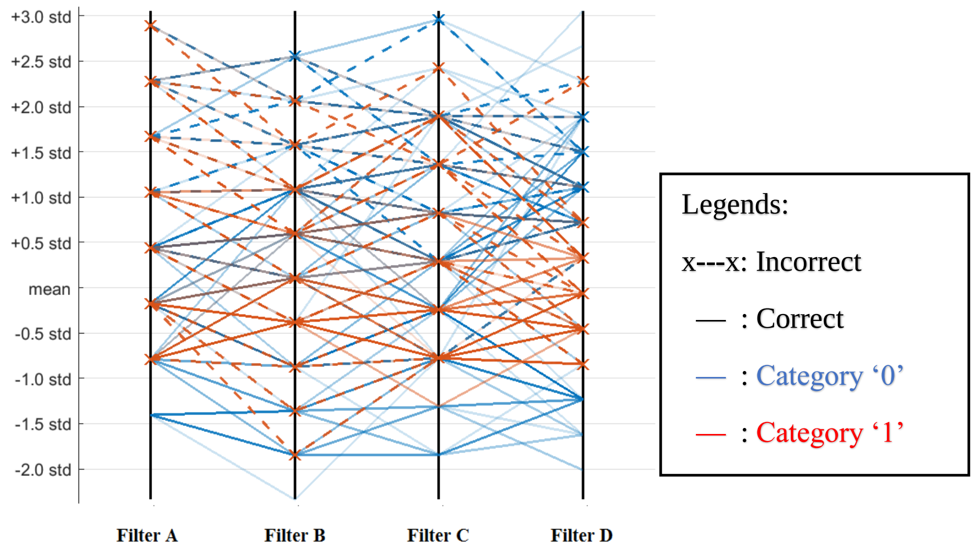

Figure 8: Parallel Coordinates Plot using Simple Tree

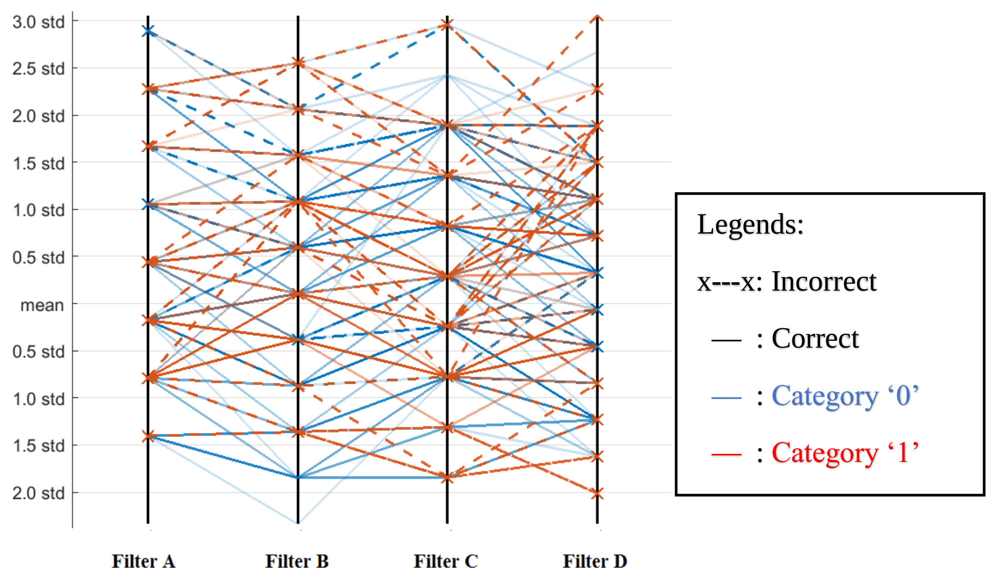

Figure 9: Parallel Coordinates Plot using Logistic Regression

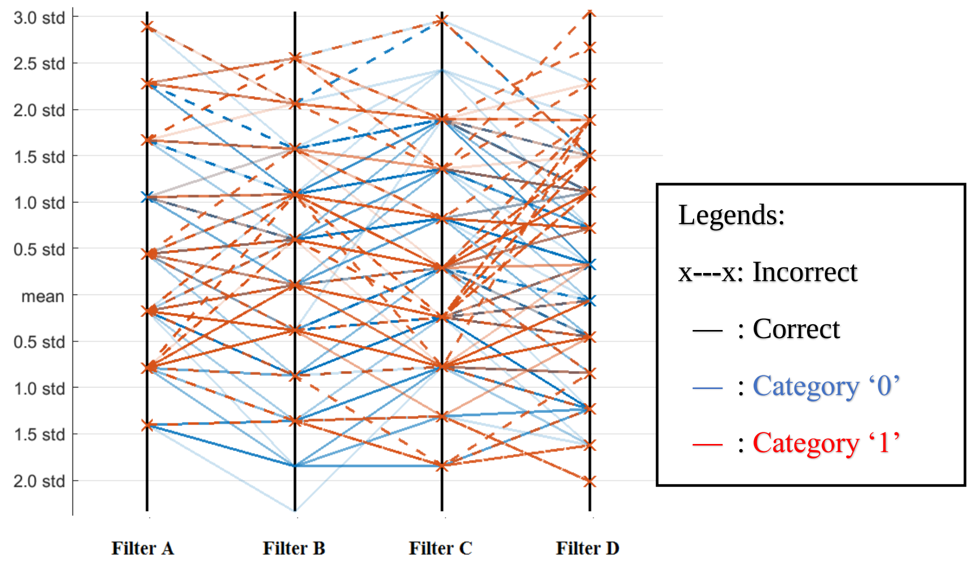

Figure 10: Parallel Coordinates Plot using Linear SVM

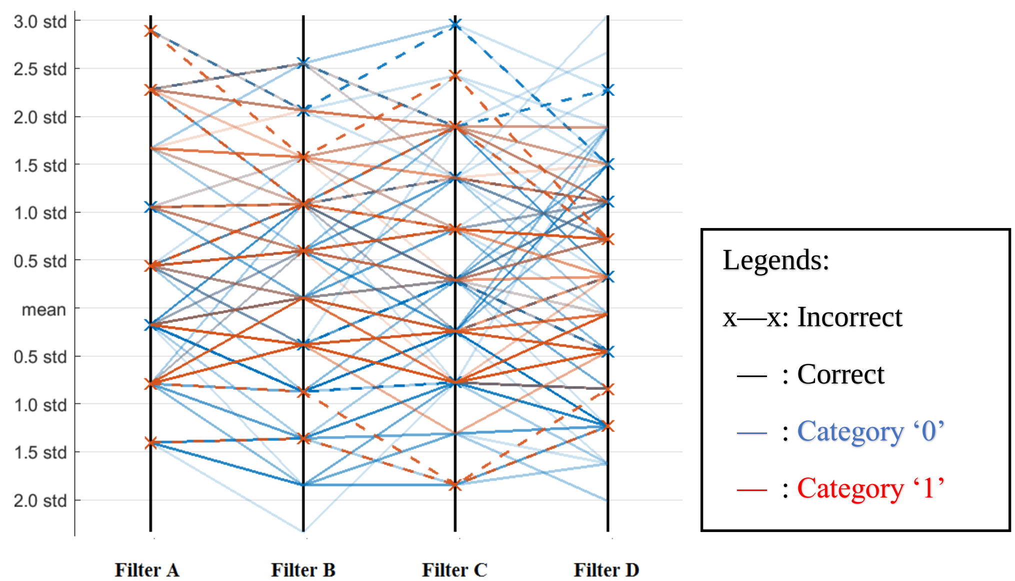

Figure 11: Parallel Coordinates Plot using Quadratic SVM

4.2 Evaluation Summary

By using the four filters we are actually trying to detect the dip and humps of the "W" pattern mentioned in the previous chapter. According to the wavelength ranges of the filters (see *Figure 5*), if light is reflected from skin, Filter B should record a relatively higher (but not too much high) value than Filter A, Filter C should record a relatively low (but not too low) value than Filter B and Filter D should record a relatively high value than Filter C (see *Figure 4*). If we observe the *Parallel Coordinates* plots for the models we used, we can clearly see that the correct category '1' lines are the ones that exhibit the above described nature. By comparing the accuracy, precision, and recall values for the models (see *Table 3*) we see that *Logistic Regression* and *Linear SVM* models exhibited lower accuracy as compared to the other two models, and their prediction speeds are also poor. However, *Simple Tree* model offered an accuracy of 83.7% and a prediction speed of 17000 observations per second, and *Quadratic SVM* had an accuracy of 94.3% and a prediction rate of 19000 observations per second. The last model seems to be the best since its accuracy and prediction speed is ahead of others. However, quadratic fitting for such a small training set clearly overfits data, therefore, the best option is the Simple Tree model.

CHAPTER 5: DISCUSSION AND CONCLUSION

The main contributions of this work are both on theoretical and practical sides. First, the use of vision sensor for skin detection is innovative since it deals with the detection at hardware level and the fact that it is a low-powered sensor makes it more interesting since it can add an active layer of privacy over the regular mobile camera. Secondly, the accuracy of skin detection I got is decent and it will encourage combined approaches to extract and detect more features from nude imagery in the future. A big database should be generated to ensure good accuracy and to allow evaluation for individual skin tones too, so that we can assess which skin tones are detected more accurately than others. If we were working with digitized image, the pixel resolution would have made a difference according to sampling theory, for example, a lower resolution could have caused the “W” pattern of skin reflectance to lose some of its peculiarity. However, I have an advantage that I am working in continuous domain and hence the filters act like analog samples of the reflected spectrum. If number of filters is increased and so does the design for each filter is improved, more accuracy can be achieved and dark skin can also be detected since improved filter might detect some hidden patterns for darker skin (may be in logarithmic scale of wavelength). The false positive cases are less because the “W” pattern for skin is very different from any other skin-like colored-object and the model trained is based on the four filters which are based on the “W” pattern. However, this does not guarantee that teens cannot find ways to get around it. Since the database is small and does not contain unique cases like nude picture in red light or underwater nude selfie, etc., the skin detector I implemented might not perform at all under such conditions.

5.1 Limitations

To be clear, the prototype implemented does not differentiate between a nude scene and a skin dominant non-nude scene. The work was constrained because I was working with minimal hardware (i.e., vision sensor chip without a lens). This means that the pixel array gets the same value and hence the image is fully-defocused. Although the composition of colors reaching the sensor can be extracted by using light filters (which I already implemented), shape or location specific features can still not be extracted without a lens. The module can detect skin dominant scenes which is an achievement towards the goal of detecting nudity considering skin is one of the most important feature in the detection of nudity, as also highlighted by Rigan Ap-apid [64]. To detect nudity, we need to use a lens in a way to make the capture partially-defocused but at the same time we need to maintain privacy. An ideal compromise between maintaining privacy and detecting skin with high accuracy would happen when the data that is captured on the pixel array is defocused enough not to be converted back to a recognizable nude image and at the same time has enough distinct features to detect nudity or high skin content.

The biggest constraint in the way to achieve a high accuracy of detection is the lack of already available dataset that can be used to sharpen our classifier model. It is difficult to capture 'real' nude scenes through the sensor to form a big database. One way could be to run a slideshow of nude images on a monitor screen and place the vision sensor in front of the screen to capture all the samples found in existing nudity related databases on the internet. However, this scheme does not work. The intrinsic glare of the monitor screen barges into data being sensed by the sensor and hence the classifier that will be learned on that data will not work for real nude scenes. Second big constraint is related to the mechanics of the sensor. Since I have introduced four light filters to extract low level features and build a classifier, I must have four values as input to the classifier every time I attempt to detect nudity. This means that I have to use four different filters to capture the same scene at a given time. Switching between the filters mechanically in front of the sensor is not practically feasible and requires time, space and mechanical power. One solution could be to use four separate sensors with different filters. This saves time though it still consumes space and electrical power. Another solution that is more feasible, in our case (fully-defocused capture i.e., one numeric value per filter), is to partition the pixel array on sensor into 4 quadrants and read four values at a time. This saves time, space, and power though it will only work for very defocused or fully-defocused captures.

5.2 Future Research Directions

In future, this work can be carried forward. The Human-Computer-Interaction part can be improved by developing the proposed mobile application. I propose that when this prototype will be implemented on a mobile device using its processor, the information of detection (either binary information i.e., detected or not detected, or multiple levels of detection) can then be sent through the background service that is operating the sensor to an application in the mobile that controls the regular camera of the device. If the accuracy of feature detection in the defocused images increases in future, nudity can also be detected while preserving the privacy and the information can be sent to an application inside that mobile device which can then control the camera according to different mitigation levels. In essence, this application will receive the information of detection of nudity and take appropriate action according to the criteria preset by the guardians/parents of the user of that mobile device. The application will at least offer three types of actions (mitigation approaches): 1) On a very basic level, it records the number of times the user (teen) has made an attempt to take a nude picture, 2) The app shuts off the device's camera for 60 seconds, and 3) The app warns the user up to a certain number of instances of attempts and then goes on to blocking the camera. Such an approach might take us one step closer to a solution to help teens make better online choices in the following ways:

- This solution caters for the age factor in a way that it presumes that the mobile application will be installed on a mobile device that belongs to an underage user.
- The processing power, time and space for the skin detection process through this sensor can easily be handled by a mobile processor. Although, processing low-level features is a compromise on accuracy of classification, detection of nudity eventually in real time on a portable device like a cell phone is the central goal and so this compromise is necessary.
- The sensor used is a low powered sensor with a set of 4 filters that can operate 24/7 and captures low level features that are easy to process in real time. This makes it compatible with real-time

applications like Skype, Facetime etc. Also, such a sensor offers privacy, since these captures by the vision sensors are defocused and are unrecognizable to any human. This means that nudity can be detected before digital capture (via regular camera) of the scene.

This future work will help fill the gap in the literature that was highlighted earlier. Though the prototype I have implemented detects only skin, the use of lens with a low-powered sensor in a mobile device may address the problem of detecting nudity in the future. Detecting nudity at this level will give more control, for instance, over capturing of an image through the regular mobile camera. Furthermore, the presence of a mobile application or background service to manage risk mitigation strategies once nudity is caught will make the solution even more comprehensive and manageable. The merging of the fields of Computer Engineering, Psychology, and Human-Computer Interaction will enable us to create and evaluate a novel solution to the problem of adolescent sexting behaviors. By de-coupling risk detection (i.e., sensor) from risk mitigation strategies employed (i.e., application), we can move toward building an integrated solution that is both technically feasible and that practically addresses the human context of the problem at hand. Finally, conducting user studies prior to developing our solution and assessing the viability of our solution after it has been built will ensure that our work translates into broader societal impacts.

However, from a more technical perspective, the accuracy of this solution needs to be improved before it can be implemented commercially. Number of filters can be increased to capture minute details in human skin reflectance pattern. Instead of using totally defocused image with one value per filter, a lens can be used to focus different light intensities onto different cells of pixel array. This means we have more information about the scene that is captured by the camera and hence more features could be extracted.

5.3 Conclusion

In this thesis, I made a skin detector module that uses a vision sensor and 4 filters to detect a skin dominant scene. As soon as the electromagnetic spectrum falls on to the sensor, the detector tells whether the reflectance pattern was skin dominant or not with an accuracy of 83%. Its novelty is in the fact that this detector does not operate on an already digitized image but detects skin at the hardware level. This ensures privacy. My main contribution is in designing the light filters for sensor to increase the accuracy of skin detection. I designed a set of four filters that detects a skewed "W" like pattern of skin reflectance that is common across different skin tones. Ten people with varying skin tones participated in creating a dataset associated with real scenes, including 150 non-skin dominant setting (e.g., a landscape or picture of an object) and 150 skin-dominant settings (e.g., an almost nude "selfie" taken in a bathroom or a close-up of one's face). The output of the sensor was in the form of four integer values, one for each filter. These form the feature vectors for learning. I trained Simple Tree, Logistic Regression, Linear SVM and Quadratic SVM on the data and test these models on 50 observations to determine which model gives the best accuracy of skin-dominancy detection with low data overfitting. Simple Tree offered the best result with an accuracy of 83.7%, precision value of 0.83 and recall value of 0.85. The long-term goal to be achieved in the future is to design a low-powered vision sensor that can be mounted on a digital camera lens on a teen's mobile device to detect and/or prevent the capture of nude imagery.

APPENDIX: HARWARE USED AND SPEC ON THE EQUIPMENT

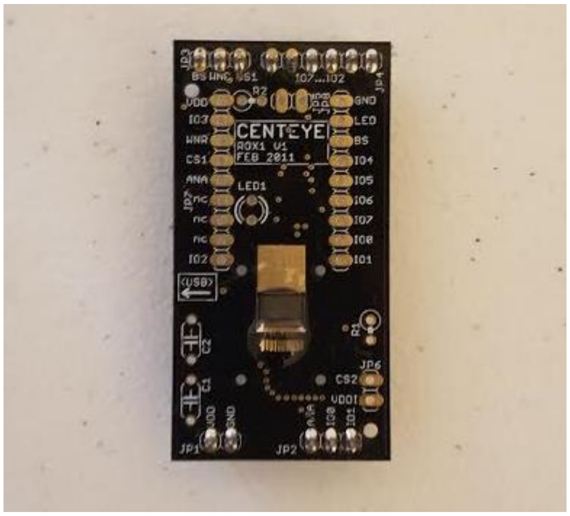

Figure 12: FireFly Sensor

Figure 13: Swatchbook (Filters)

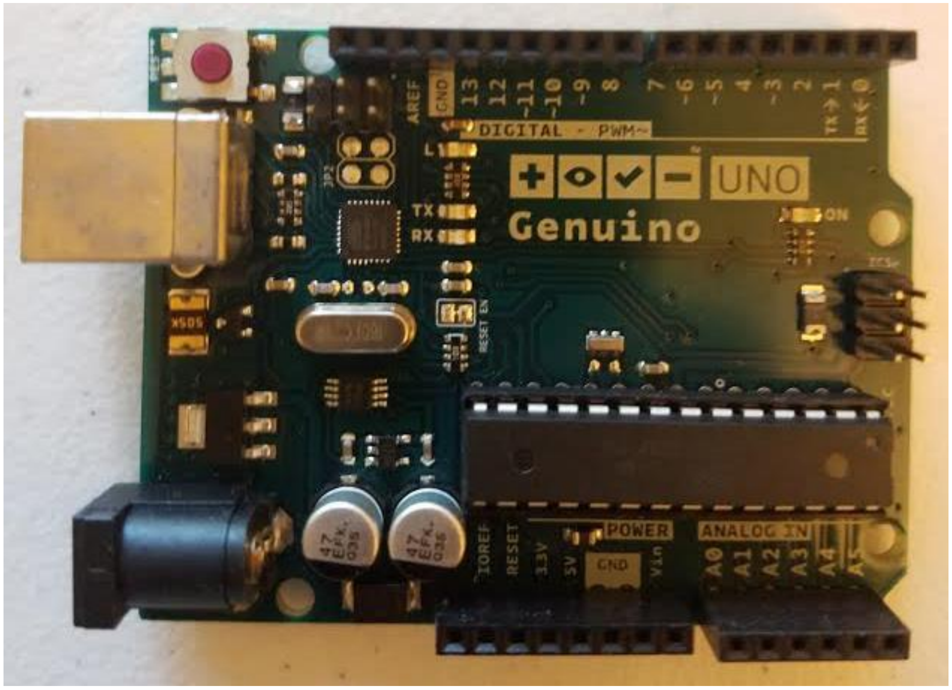

Figure 14: Arduino/Genuino Board

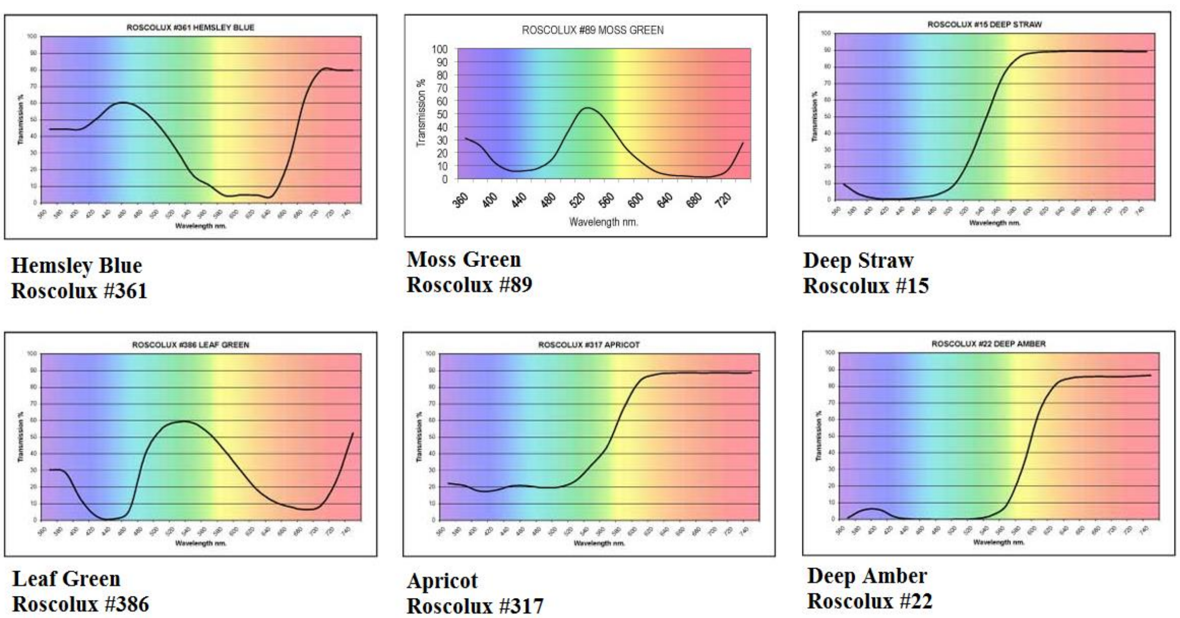

Figure 15: SEDs of some common filters

LIST OF REFERENCES

01 - Bryce G. Westlake - Liking and hyperlinking Community detection in online child sexual exploitation networks.pdf, n.d.

02 - Mofakharul Islam - Real-time detection of children’s skin on social networking sites using Markov random field modelling.pdf, n.d.

03 - K K Bhoyar - Skin Color Detection Model Using Neural Networks and its Performance Evaluation.pdf, n.d.

04 - Shubhangi D Raverkar - Classification of YouTube Metadata using Shark Algorithm.pdf, n.d.

05 - S Wijesinghe - Parental Control and Filtering System.pdf, n.d.

06 - Pedro Ivan Tello Flores - Approach of RSOR Algorithm Using HSV Color Model for Nude Detection in Digital Images.pdf, n.d.

07 - Mateus de Castro Polastro - NuDetective a Forensic Tool to Help Combat Child Pornography through Automatic Nudity Detection.pdf, n.d.

08 - Clayton Santos - NUDITY DETECTION BASED ON IMAGE ZONING.pdf, n.d.

09 - Mateus de Castro Polastro - A statistical approach for identifying videos of child pornography at crime scenes.pdf, n.d.

10 - Thomas Vanhove - Towards the Design of a Platform for Abuse Detection in OSNs using Multimedial Data Analysis.pdf, n.d.

11 - Chirag Ahuja - Detection of Nude Images on Large Scale Using Hadoo.pdf, n.d.

12 - Jafar Shayan - An overview of objectionable image detection.pdf, n.d.

13 - Medeni Soysal - Multimodal concept detection in broadcast media KavTan.pdf, n.d.

14 - Alierza Behrad - Content-based obscene video recognition by combining 3D spatiotemporal and motion-based features.pdf, n.d.

15 - Nicole A. Poltash - SNAPCHAT AND SEXTING A SNAPSHOT OF BARING YOUR BARE ESSENTIALS.pdf, n.d.

16 - Awais Adnan - RGB and Hue Color in Pornography Detection.pdf, n.d.

17 - M. Grega - INACT — INDECT Advanced Image Cataloguing Tool.pdf, n.d.

18 - Digambar Povar - Digital Image Evidence Detection Based on Skin Tone Filtering Technique.pdf, n.d.

19 - Luca Giangiuseppe Esposito - A Multiple Classifier Approach for Detecting Naked Human Bodies in Images.pdf, n.d.

20 - Hakan Sevimli - Adult Image Content Classification Using Global Features and Skin Region Detection.pdf, n.d.

21 - Nilesh J.Uke - Detecting Pornography on Web to Prevent Child Abuse – A Computer Vision Approach.pdf, n.d.

22 - Pedro Monteiro da Silva Eleuterio - An overview of NuDetective Forensic Tool and its usage to combat child pornography in Brazil.pdf, n.d.

23 - Yong Lin - Machine Learning Application On Detecting Nudity In Images.pdf, n.d.

24 - Lorenz Tombo Ras - NUDITY PROHIBITION SYSTEM.pdf, n.d.

25 - Shubhangi T Patil - Image Cataloguing Tool using Descriptor for Forensic Application.pdf, n.d.

26 - Christian Platzer - Skin Sheriff A Machine Learning Solution for Detecting Explicit Images.pdf, n.d.

27 - Fernando Roberti de Siqueira - Adaptive Detection of Human Skin in Color Images.pdf, n.d.

28 - Supreet Sidhu - CONTENT BASED IMAGE RETRIEVAL A REVIEW.pdf, n.d.

29 - Pedro Monteiro da Silva Eleuterio - Identification of High-Resolution Images of Child and Adolescent Pornography at Crime Scenes.pdf, n.d.

30 - Mateus de Castro Polastro - uick Identification of Child Pornography in Digital Videos.pdf, n.d.

31 - Satrio Dewantono - Development of A Real-Time Nudity Censorship System on Images.pdf, n.d.

32 - Wayne Kelly - Screening for Objectionable Images A Review of Skin Detection Techniques.pdf, n.d.

33 - Thomas Deselaers - Bag-of-Visual-Words Models for Adult Image Classification and Filtering.pdf, n.d.

34 - Rainer Lienhart - FILTERING ADULT IMAGE CONTENT WITH TOPIC MODELS.pdf, n.d.

35 - Xinlu Wang - Nude image detection based on SVM.pdf, n.d.

36 - Ali Selamat - Fuzzy Mamdani Inference System Skin Detection.pdf, n.d.

37 - Xiaoyin Wang - An Adult Image Recognizing Algorithm Based on Naked Body Detection.pdf, n.d.

38 - Vijayendar G - Integrated Approach to Block Adult Images in Websites.pdf, n.d.

39 - Ana Paula B. Lopes - Nude Detection in Video using Bag-of-Visual-Features.pdf, n.d.

40 - Giuseppe Amato - Detection of Images with Adult Content for Parental Control on Mobile Devices.pdf, n.d.

41 - Mr Alisdair A. Gillespie - Adolescents accessing indecent images of children.pdf, n.d.

42 - Glen Thompson - Automatic detection of child pornography.pdf, n.d.

43 - Rigan Ap-apid - Image-Based Pornography Detection.pdf, n.d.

44 - Rehanullah Khan - An Adaptive Multiple Model Approach for Fast Content-Based Skin Detection in On-Line Videos.pdf, n.d.

45 - Bei-bei Liu - Pornographic Images Detection Based on CBIR and Skin Analysis.pdf, n.d.

46 - CITIZEN'S GUIDE TO U.S. FEDERAL LAW ON CHILD PORNOGRAPHY [WWW Document], n.d. URL https://www.justice.gov/criminal-ceos/citizens-guide-us-federal-law-child-pornography

47 - Teens and Sexting [WWW Document], n.d. URL http://www.ncdsv.org/images/pewinternet_teensandsexting_12-2009.pdf

48 - Nielson Colton - Review of digital image security in Dermatology, n.d.

49 - PhotoDNA Lets Google, FB and Others Hunt Down Child Pornography Without Looking at Your Photos [WWW Document], n.d. URL https://petapixel.com/2014/08/08/photodna-lets-google-facebook-others-hunt-down-child-pornography-without-looking-at-your-photos/

50 - Fighting the harmful content problem [WWW Document], n.d. URL https://www.microsoft.com/en-us/photodna

51 - Apple toughens iCloud security after celebrity breach [WWW Document], n.d. URL http://www.bbc.com/news/technology-29237469

52 - Parental control software [WWW Document], n.d. URL http://www.myoutofcontrolteen.com/mr-PCS.html

53 - Sook-Jung Lee and Young-Gil Chae. 2012. Balancing participation and risks in children's Internet use: the role of internet literacy and parental mediation. Cyberpsychology, Behavior and Social Networking 15, 5: 257–262 [WWW Document], n.d. URL http://online.liebertpub.com/doi/abs/10.1089/cyber.2011.0552

54 - Sahara Byrne - Peers, Predators, and Porn: Predicting Parental Underestimation of Children's Risky Online Experiences [WWW Document], n.d. URL http://onlinelibrary.wiley.com/doi/10.1111/jcc4.12040/abstract

55 - Eva Lievens - Bullying and sexting in social networks: protecting minors from criminal acts or empowering minors to cope with risky behaviour? [WWW Document], n.d. URL https://biblio.ugent.be/publication/6953092

56 - Kovac - Human Skin Colour Clustering for Face Detection.pdf, n.d.

57 - Yu-Chun LinPornography Detection Using Support Vector Machine.pdf, n.d.

58 - Snapchat sexting and the predators of Kik: The apps your children need to stay away from [WWW Document], n.d. URL http://www.dailytelegraph.com.au/news/snapchat-sexting-and-the-predators-of-kik-the-apps-your-children-need-to-stay-away-from/news-story/ad63dc7962ff4098b0785d2671f5c89c

59 - 9 apps to improve your sext life [WWW Document], n.d. URL http://www.nydailynews.com/news/national/9-apps-improve-sext-life-article-1.2465165

60 - Skin Color Detection Model Using Neural Networks and its Peformance Evaluation.pdf, n.d.

61 - Redmond, Teresa (22 February 2010). "Ringwood community addresses sexting". NorthJersey.com. North Jersey Media Group. Retrieved 30 May 2010.

62 - Sexting: A new, digital vehicle for intimate partner aggression? Michelle Drouin http://www.sciencedirect.com/science/article/pii/S0747563215002836

63 - Elli Angelopoulou, "The Reflectance Spectrum of Human Skin",. December 1999 http://repository.upenn.edu/cgi/viewcontent.cgi?article=1616&context=cis_reports

64 - Alex Smola and S.V.N. Vishwanathan, “Introduction to Machine Learning” http://alex.smola.org/drafts/thebook.pdf

65 - Duan Lijuan, Cui Guoqin, “A Hierarchical Method for Nude Image Filtering” http://en.cnki.com.cn/Article_en/CJFDTOTAL-JSJF200205004.htm