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

Muhammad Uzair Tariq
University of Central Florida
Orlando, USA
uzairtq@gmail.com

Arup Kumar Ghosh
University of Central Florida
Orlando, USA
arupkumar.ghosh@ucf.edu

Karla Badillo-Urquiola
University of Central Florida
Orlando, USA
kcurquiola10@knights.ucf.edu

Abhiditya Jha
University of Central Florida
Orlando, USA
abhiditya@knights.ucf.edu

Sanjeev Koppal
University of Florida
Gainsville, USA
sjkoppal@ece.ufl.edu

Pamela J. Wisniewski
University of Central Florida
Orlando, USA
pamwis@ucf.edu

***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 on high fidelity images 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. Our research takes a first step in trying to accomplish this goal: In this paper, we 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), we were able to accurately detect a skin dominant scene with 83.7% accuracy, 83% precision, and 85% recall. Our long-term goal 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, we discuss the limitations of this work toward this larger goal, as well as future research directions.**

***Keywords—filters, skin detection, nudity, low-powered sensor, machine learning***

## I. Introduction

The advent of mobile smart devices, digital image capture, and multi-media messaging services has created a phenomenon known as "sexting [1]," 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 [2]. 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 by prior research [1], [3]. To prevent such 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 after 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 could 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 would possibly be transmitted to a server for additional processing, already negating 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.

As part of this work, we 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. From this, 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). We 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

978-1-5386-6133-8/18/$31.00 ©2018 IEEE

detects and/or prevents the digital capture of nude imagery through a regular camera.

## II. Background

In the following section, we provide an overview of the current literature regarding adolescent mobile risks and nudity detection approaches.

### *A. Risks Associated with Digitized Nudity*

A recent study found 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 [4]. Teen sexting behaviors are perpetuated by mobile technologies and by several direct messaging applications that are available on smartphones like Kik, Snapchat, and AskFM [5]. Unfortunately, such activities often fall under the jurisdiction of child pornography laws. Child pornography is illegal, and the federal law states that "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" [2].

### *B. The Current State of the Art in Nudity 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 [6] developed a method for detecting adult nudity in videos based on a bag-of-visual-features representation for frames; Kovac et al. [7] 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. [8] used a Support VectorMachine (SVM), which has learning skills in the image detection of human nudity. Amato [9] 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 [10]. PhotoDNA [11], 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. We 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 [12]–[20], 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.

In the following section, we explain in detail the implementation of the skin detector. First, we go into the detail of hardware including the specific design of filters. Then, we explain the setting of the experiment that we performed to test the accuracy of the skin detector. In the results sections, we present and interpret the results of the experiment. Finally, we discuss the implications of this work, provide the limitations associated with the implementation, and outline 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.

### *C. Skin Detection as an Integral Part of 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 is skin [21], [22]. Islam et al. in [23] states:

> "Nudity and pornography have a direct link with human skin. In fact, no pornography can exist without exposure of human skin… A wide range of image processing applications exist, where skin detection is playing a crucial role. Using colour as a detection cue has long being recognized as a robust feature and has become a popular choice in human skin detection techniques."

With this in mind, our work draws from Angelopoulou et al. [24], which showed that skin reflectance exhibits a "W" pattern in wavelength domain. Therefore, we design light filters to detect this pattern through the low-powered vision sensor to improve skin detection on low fidelity imagery.

## III. Hardware Configuration and Components

This section describes the physical hardware components and configuration of our prototype. Since the basic vision sensors (e.g., FireFly Vision Sensor) cannot be connected to a cell phone without dismantling the mobile device and modifying its hardware. Therefore, for the sake of this proof-of-concept experiment, we used 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). These sensors, in theory, can be connected to existing mobile processors once our proof-of-concept has matured into a viable solution for the problem.

### *A. Hardware Components*

*Processor*: We used Arduino/Genuino Board using the open-sourceArduino Software (IDE) which makes it 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)*: We used the FireFlyBig [25]. 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. We used the former one.

### B. *Hardware Configuration*

The sensor is connected to a processor to send and receive signals. For our purposes, we did not need the 8-bit port. The analog output of the FireFly chip is providing us with pixel values of the pixel array, serially. Figure 1 shows an actual setup of our hardware prototype.

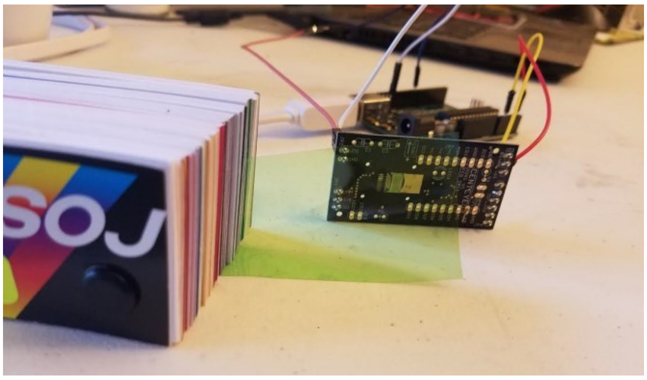

Figure 1. Setup of the Hardware Prototype

Process for configuring the hardware:

- 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 to put between the sensor and the subject. Check to see if the value received changes with different filters.

These steps ensured that everything was connected and working properly prior to running the experiment.

## IV. DESIGNING CUSTOMIZED LIGHT FILTERS

Before running the experiment, we needed to understand what data is received from the sensor and how filters work on the sensor to provide the useful data that we need. We designed some specific filters to meet these requirements, using basic and commercially available filters.

### A. *Sensors, Sensor Data, and Light Filters*

In the hardware configuration, we intentionally used a sensor with no lens. Therefore, the data was 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, our goal was to defocus the image in order to maintain privacy. Defocusing the image means less feature detection (i.e., low fidelity imagery), which is a key trade-off (between preserving privacy and accurate nudity detection) explicitly made by this research design. This also meant that we 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), we took four different captures of the same object/scene with four different light-filters each. We designed these four filters.

### B. *Filter Design*

The selection of the four filters (Filters A, B, C, and D) that we designed can be justified by Angelopoulou et al. [24]. 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). It shows the pattern of reflectances of various tones of human skin for wavelengths of visible light. The main goal of designing filters is to detect the approximate position of the two peculiar dips and a hump shown in the "W" pattern.

Ideally, the filters that we select should act as band pass filters allowing four roughly-distinct range of wavelengths to pass through to the sensor. Filter A should transmit all the wavelengths from 500 to 550nm. Filter B should be designed to transmit only the wavelengths from 550 to 565nm. Filter C should pass wavelengths ranging from 565nm to 580nm and Filter D should allow all the wavelengths above 580 to pass through. The ideal filter will pass the required range of wavelength with 100 percent transmission rate and stop the rest of the wavelengths completely. We 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, we discuss how the SED curve of some commonly available filters can be used to design a required filter. Deep Straw (Roscolux #15) and Leaf Green (Roscolux #386) filters were combined to make Filter B.

Similarly, for Filter 'A' we used Hemsley Blue and Moss Green, for Filter 'C' we used Apricot and Leaf Green, and for Filter 'D' we just used single filter i.e., Deep Amber. 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. They can be improved given more variety of filters available in market.

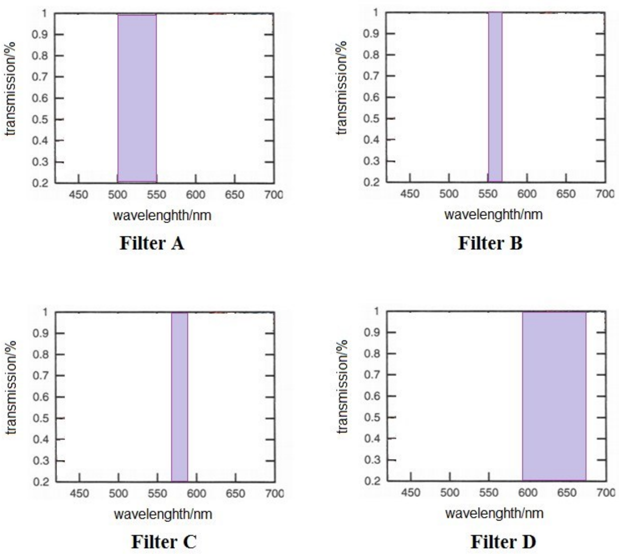


Figure 2. Ideal Light Filters Needed for the Experiment

## V. Study Design

### A. Data Collection

We used the FireFly Vision chip [25] 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. 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.

Ten individuals, including the first author, 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 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. We did not include very dark-skinned people (black), since the "W" pattern for the reflectance of spectrum mentioned by

| 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<br>Split criterion: Gini's diversity index<br>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<br>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<br>Multiclass method: One vs One |

Table 1: Model Results

Angelopoulou et al. [24] does not exist for them.

*B. Analysis and Classification*

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 we 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. we chose the above-mentioned algorithms based on Alex Smola et al. [21]. Since, we are detecting skin through only four filters, low accuracy is already expected and therefore achieving a classification accuracy of 80% should be considered a success. Skin detectors proposed in [23], [26]–[32] got accuracy around 70% to 90%.

To perform the analysis, we 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. We tested four models based on the algorithms selected above and train the models with 250 observations (125 being skin-dominant). Then we tested the models on 50 observations (25 being skin-dominant) to record the accuracy, precision and recall of the models. Next, we provide results from the four models.

## VI. RESULTS

The results of the four models are shown in Table 1 with figures that illustrate these of these models shown in Figures 3-6. To do the analysis, we 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, we applied the following four classification models from the options: Simple Tree, Logistic Regression, Linear SVM and Quadratic SVM. We 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. We 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.

The four plots shown in Figures 3-6 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).

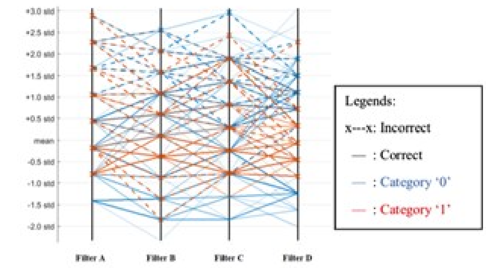


Figure 3. Parallel Coordinates Plot using Simple Tree

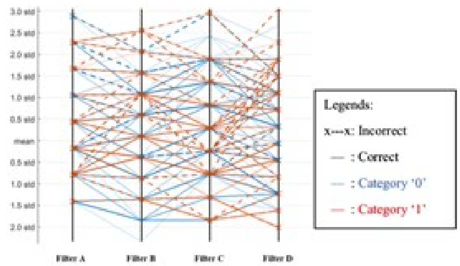


Figure 4. Parallel Coordinates Plot using Logistic Regression

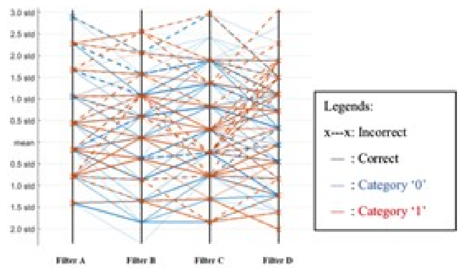


Figure 5. Parallel Coordinates Plot using Linear SVM

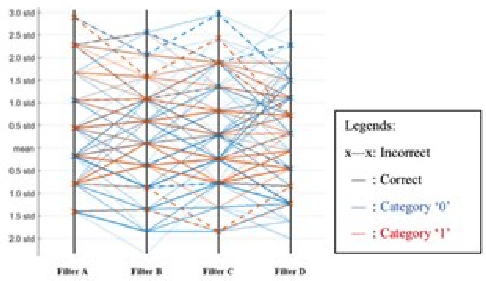


Figure 6. Parallel Coordinates Plot using Quadratic SVM

### *A. Evaluation Summary*

By using the four filters we aimed to detect the dip and humps of the "W" pattern mentioned earlier. According to the wavelength ranges of the filters (see Fig. 2), 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 2). 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 1) 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.

## VII. Discussion

The main contributions of this work are both on theoretical and practical sides. First, the use of a 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 the skin detection will encourage combined approaches to extract and detect more features from nude imagery in the future. A large 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 images, 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, the advantage of this research is that we are working in with a continuous domain and hence the filters act like analog samples of the reflected spectrum.

If the number of filters is increased, the design for each filter is improved. More accuracy can be achieved and dark skin can also be detected since the 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. The skin detector may not perform at all under unique cases like nude picture in red light or underwater nude selfie, because the database is small.

### *A. Limitations*

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. To be clear, the prototype implemented does not differentiate between a nude scene and a skin dominant non-nude scene. The work was constrained due to 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 (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 features in the detection of nudity, as also highlighted by Rigan Ap-apid [21]. 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 to achieving a high accuracy of detection is the lack of already available datasets 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.

Another constraint is related to the mechanics of the sensor. Since we have introduced four light filters to extract low level features and build a classifier, we must have four values as inputs to the classifier every time we attempt to detect nudity. This means that four different filters must be used to capture the same scene at a given time. Switching between the filters mechanically in front of the sensor is not 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 the 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.

Finally, we are not providing a hardware-ready solution. Instead, we are presenting a proof-of-concept as an implementable idea that mobile companies can take into account while designing their products.

### *B. Future Research*

Future research will need to test this approach on mobile devices. Using mobile processors, the information of detection (either binary information i.e., detected or not detected, or multiple levels of detection) could 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, 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 actions according to the criteria preset by the guardians/parents of the user of that mobile device. The application might offer multiple types of mitigation approaches, including: 1) On a very basic level, it might record the number of times the user (teen) has tried to take a nude picture, 2) The app could shut off the device's camera for 60 seconds, and 3) The app may warn the user up to a certain number of instances of attempts and then goes on to blocking the camera. Depending on the balance between parental control and teen self-regulation appropriate for a given parent-child relationship [33]–[35], these parameters may be shifted to adjust to the needs of the users. Such approaches might take us one step closer to solutions that help teens make better online choices. Some reasons why we believe our approach presents a feasible solution moving forward include:

- 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. Therefore, the skin detection algorithms do not need to differentiate between adults versus minors.
- 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.

Although the current prototype detects only skin, the use of lens with a low-powered sensor in a mobile device may be able to 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.

However, the sensor would need to be integrated with the camera lens by phone manufacturers or done after-market in a way that ensures that it cannot easily be removed from the mobile device by teens trying to circumvent its functionality. Thus, to move forward with a consumer-ready product, multiple stakeholders (e.g., phone manufacturers, operating system providers, designers, engineers, and researchers) need to work together to develop an end-to-end solution that both detects risky sexting behaviors and works to mitigate these behaviors before they endanger the safety of our youth. 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 [36]. 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 impact.

## VIII. Conclusion

We developed 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 onto 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. Our main contribution is in designing the light filters for sensor to increase the accuracy of skin detection. 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 (or any minor under the age of eighteen) mobile device to detect and/or prevent the capture of nude imagery.

## References


[1] "Sexting," *Wikipedia*. 18-Feb-2018.

[2] "Citizen's Guide To U.S. Federal Law On Child Pornography | CRIMINAL-CEOS | Department of Justice." [Online]. Available: https://www.justice.gov/criminal-ceos/citizens-guide-us-federal-law-child-pornography. [Accessed: 05-Feb-2018].

[3] P. Wisniewski, H. Xu, M. B. Rosson, D. F. Perkins, and J. M. Carroll, "Dear Diary: Teens Reflect on Their Weekly Online Risk Experiences," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2016, pp. 3919–3930.

[4] A. Lenhart, "Teens and sexting," *Pew internet & American life project*, vol. 1, pp. 1–26, 2009.

[5] "Apps your kids should not use," 06-Jun-2014. [Online]. Available: https://www.dailytelegraph.com.au/news/snapchat-sexting-and-the-predators-of-kik-the-apps-your-children-need-to-stay-away-from/news-story/ad63dc7962ff4098b0785d2671f5c89c. [Accessed: 05-Feb-2018].

[6] T. Deselaers, L. Pimenidis, and H. Ney, "Bag-of-visual-words models for adult image classification and filtering," in *2008 19th International Conference on Pattern Recognition*, 2008, pp. 1–4.

[7] J. Kovac, P. Peer, and F. Solina, "Human skin color clustering for face detection," in *The IEEE Region 8 EUROCON 2003. Computer as a Tool.*, Ljubljana, Slovenia, Slovenia, 2003, vol. 2, pp. 144–148 vol.2.

[8] Y.-C. Lin, H.-W. Tseng, and C.-S. Fuh, "Pornography detection using support vector machine," *16th IPPR Conference on Computer Vision, Graphics and Image Processing (CVGIP 2003)*, Jan. 2003.

[9] G. Amato, P. Bolettieri, G. Costa, F. la Torre, and F. Martinelli, "Detection of Images with Adult Content for Parental Control on Mobile Devices?," in *Proceedings of the 6th International Conference on Mobile Technology, Application & Systems*, New York, NY, USA, 2009, pp. 35:1–35:5.

[10] G. Burgett, "PhotoDNA Lets Google, FB and Others Hunt Down Child Pornography Without Looking at Your Photos," 08-Aug-2014. [Online]. Available: https://petapixel.com/2014/08/08/photodna-lets-google-facebook-others-hunt-down-child-pornography-without-looking-at-your-photos/. [Accessed: 05-Feb-2018].

[11] "Fighting the harmful content problem." [Online]. Available: https://www.microsoft.com/en-us/photodna. [Accessed: 05-Feb-2018].

[12] S. D. Raverkar and M. Nagori, "Classification of YouTube Metadata using Shark Algorithm," *International Journal of Computer Applications*, vol. 132, no. 9, pp. 18–21, Dec. 2015.

[13] T. Vanhove, P. Leroux, T. Wauters, and F. D. Turck, "Towards the design of a platform for abuse detection in OSNs using multimedial data analysis," in *2013 IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*, 2013, pp. 1195–1198.

[14] A. Behrad, M. Salehpour, M. Ghaderian, M. Saiedi, and M. N. Barati, "Content-based obscene video recognition by combining 3D spatiotemporal and motion-based features," *J Image Video Proc*, vol. 2012, no. 1, p. 23, Dec. 2012.

[15] A. Adnan and M. Nawaz, "RGB and Hue Color in Pornography Detection," in *Information Technology: New Generations*, Springer, Cham, 2016, pp. 1041–1050.

[16] N. Uke, C. Ravindra, and R. Thool, "Detecting Pornography on Web to Prevent Child Abuse – A Computer Vision Approach," *International Journal of Scientific & Engineering Research,* vol. 3, pp. 1–3, Apr. 2013.

[17] P. M. da Silva and M. de C. Polastro, "An overview of NuDetective Forensic Tool and its usage to combat child pornography in Brazil," presented at the XLIII Jornadas Argentinas de Informática e Investigación Operativa (43JAIIO)-VI Workshop de Seguridad Informática (Buenos Aires, 2014), 2014.

[18] P. Eleuterio and M. Polastro, "Identification of High-Resolution Images of Child and Adolescent Pornography at Crime Scenes," *IJoFCS*, pp. 49–59, Jan. 2010.

[19] M. Vidigal de Castro, "Quick Identification of Child Pornography in Digital Videos - Semantic Scholar," 08-Nov-2012. [Online]. Available: https://www.semanticscholar.org/paper/Quick-Identification-of-Child-Pornography-in-Digit-Castro/07faa38d4d0e9d14d72bd049362efa83fae78ee3. [Accessed: 05-Feb-2018].

[20] A. P. B. Lopes, S. E. F. d Avila, A. N. A. Peixoto, R. S. Oliveira, M. d M. Coelho, and A. d A. Araújo, "Nude Detection in Video Using Bag-of-Visual-Features," in *2009 XXII Brazilian Symposium on Computer Graphics and Image Processing*, Rio de Janiero, Brazil, 2009, pp. 224–231.

[21] "Introduction to Machine Learning by Alexander J. Smola." [Online]. Available: https://www.goodreads.com/book/show/17876636-introduction-to-machine-learning. [Accessed: 01-Mar-2018].

[22] L. Duan, W. Gao, G. Cui, and H. Zhang, "A Hierarchical Method for Nude Image Filtering," *Journal of Computer-Aided Design and Computer Graphics*, vol. 14, Jan. 2002.

[23] M. Islam, P. A. Watters, and J. Yearwood, "Real-time detection of children's skin on social networking sites using Markov random field modelling," *Information Security Technical Report*, vol. 16, no. 2, pp. 51–58, May 2011.

[24] E. Angelopoulou, "The Reflectance Spectrum of Human Skin," *Technical Reports (CIS)*, Dec. 1999.

[25] "Centeye." [Online]. Available: http://www.centeye.com/. [Accessed: 01-Mar-2018].

[26] K. K. Bhoyar and O. G. Kakde, "Skin Color Detection Model Using Neural Networks and its Performance Evaluation," *Journal of Computer Science*, pp. 963–968, 2010.

[27] D. Povar, D. S. Vidyadharan, and K. L. Thomas, "Digital Image Evidence Detection Based on Skin Tone Filtering Technique," in *Advances in Computing and Communications*, 2011, pp. 544–551.

[28] H. Sevimli *et al.*, "Adult Image Content Classification Using Global Features and Skin Region Detection," in *Computer and Information Sciences*, Springer, Dordrecht, 2011, pp. 253–258.

[29] F. Roberti de Siqueira, "Adaptive Detection of Human Skin in Color Images - Semantic Scholar," 2013. [Online]. Available: /paper/Adaptive-Detection-of-Human-Skin-in-Color-Images-Siqueira-Schwartz/345198bfb00f80a3be7309239abfc1b87f67fcfd. [Accessed: 05-Feb-2018].

[30] W. Kelly, A. Donnellan, and D. Molloy, "Screening for Objectionable Images: A Review of Skin Detection Techniques," in *2008 International Machine Vision and Image Processing Conference*, Portrush, Ireland, 2008, pp. 151–158.

[31] R. Khan, J. Stöttinger, and M. Kampel, "An Adaptive Multiple Model Approach for Fast Content-based Skin Detection in On-line Videos," in *Proceedings of the 1st ACM Workshop on Analysis and Retrieval of Events/Actions and Workflows in Video Streams*, New York, NY, USA, 2008, pp. 89–96.

[32] B. b Liu, J. y Su, Z. m Lu, and Z. Li, "Pornographic Images Detection Based on CBIR and Skin Analysis," in *2008 Fourth International Conference on Semantics, Knowledge and Grid*, Beijing, China, 2008, pp. 487–488.

[33] P. Wisniewski, A. K. Ghosh, M. B. Rosson, H. Xu, and J. M. Carroll, "Parental Control vs. Teen Self-Regulation: Is there a middle ground for mobile online safety?," in *Proceedings of the 20th ACM Conference on Computer Supported Cooperative Work & Social Computing*, Portland, OR, USA, 2017.

[34] A. K. Ghosh, K. Badillo-Urquiola, M. B. Rosson, H. Xu, J. Carroll, and P. Wisniewski, "A Matter of Control or Safety? Examining Parental Use of Technical Monitoring Apps on Teens' Mobile Devices," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Montréal, Canada, 2018.

[35] A. K. Ghosh, K. Badillo-Urquiola, S. Guha, J. LaViola Jr., and P. Wisniewski, "Safety vs. Surveillance: What Children Have to Say about Mobile Apps for Parental Control," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Montréal, Canada, 2018.

[36] A. T. Pinter, P. Wisniewski, H. Xu, M. B. Rosson, and J. M. Carroll, "Adolescent Online Safety: Moving Beyond Formative Evaluations to Designing Solutions for the Future," in *Proceedings of the International Conference on Interaction Design and Children*, 2017.