

2009

Automatic detection of child pornography

Glen Thompson
Edith Cowan University

DOI: [10.4225/75/57b2812540cc9](https://doi.org/10.4225/75/57b2812540cc9)

Originally published in the Proceedings of the 7th Australian Digital Forensics Conference, Edith Cowan University, Perth Western Australia, December 3rd 2009.

This Conference Proceeding is posted at Research Online.

<http://ro.ecu.edu.au/adf/60>

Automatic detection of child pornography

Glen Thompson
secau – Security Research Centre
School of Computer and Security Science
Edith Cowan University

Abstract

The ability to automatically identify and classify child pornography has been a problem for researchers in recent times and an issue that law enforcement agencies have faced since the beginning of digital trafficking of child pornography. In this research in progress paper the researcher outlines the importance of automating child pornography identification for law enforcement, reviews literature relevant to the field in computer vision, and establishes the key areas to which the researcher is taking this research and what direction they plan to investigate further as a PhD thesis.

Keywords

Child pornography, erotica, child abuse, computer vision, automation, digital forensics, law enforcement, internet

INTRODUCTION

Before the introduction of the internet, the availability of child pornography was reported as on the decline (Jenkins 2001). Since its emergence, however, the internet has made child pornography a much more accessible and available means of trafficking across borders (Biegel 2001; Jenkins 2001; Wells, Finkelhor et al. 2007). The internet as it is at present is made up of a vast array of protocols and networks where traffickers can anonymously share large volumes of illegal material amongst each other from locations with relaxed or non-existent laws that prohibit the possession or trafficking of illegal material. Likewise the internet is home to new developing social networks on the world wide web where young people are attracted to sharing their personal information amongst friends and family and inevitably become targets of predators.

The volume and availability of such content, or targets for predators can be an overwhelming task for law enforcement to track and/or catalogue. In general cases image collections can range in the thousands (Taylor and Quayle 2003), and to assist in the identification and classification of child pornography within these large collections, the research of this author's PhD seeks to establish a automated method of identifying and classifying material that has a high probability of being child pornography. This paper establishes the working progress of the author as they review existing relevant literature and hypothesise possible methods of identification and classification.

Significance of Research

The objective of this research is to inevitably invent a new robust method of single image interrogation (excluding multiple frames or video information) that has a low false positive and low false negative detection rate. Speed of detection is a secondary objective of this proposed research with accuracy of results being the primary objective. With such a method available to law enforcement, large quantities of image data could be filtered at faster speeds than a human being and allow investigators to focus on relevant material.

Child Pornography

In order to design an automated approach to identifying child pornography one must understand what the term 'child pornography' actually means.

A degree of uncertainty has always shrouded the term 'child pornography'. However there has been some discussion made into the topic such as that made by Taylor and Quayle (Taylor and Quayle 2003). Looking at the term itself one can claim that by 'pornography' one means images of naked people. By that example the activity of Naturism, the cultural or political movement that advocates and defends social nudity, is an example of imagery which may be classified as pornography. Such an example however is a legitimate and appropriate activity, but yet is similar in description to what may also be termed as naked images of children and adults, or erotica.

Erotica can be described as sexually explicit material of consenting adults which is non-violent, non-degrading, or pleasurable (Marshall and Barrett 1990). However, pornography is described as an activity where one person is portrayed as powerless or non-consenting, and that because children cannot give consent to any sexual activities this is what can, at

a minimum, be defined as child pornography (Taylor and Quayle 2003). In its extreme, child pornography is the portrayal of a sexual assault against a child. In either case an image of such is classified as a picture of a crime in progress and possession of such material is deemed illegal in most countries.

Understanding this makes it clear that in order to appropriately classify child pornography, child pornography requires a system which does not account just for images of naked children, but has an understanding of what child pornography actually is. As an example, the system must not only account for images of naked children being assaulted, but children in positions of exploitation or are depicted as being victims of powerless situations. In these cases even images of clothed children can be classified as victims of molestation, making traditional forms of adult pornography detection inappropriate and insufficient as a solution for child pornography.

Traditional pornography detection

Of the bulk of pornography, or skin, filtering research, the use of back-propagation artificial neural networks (ANN) is prevalent as the primary means of classifying images as skin or pornography prevalence with high efficiency and low false positive detection rate (Forsyth and Fleck 1999; Jones and Rehg 2002; Zhu, Wu et al. 2004; Zheng and Gao 2005).

The existing methods that use ANN first apply a number of ‘smoothing’ methods (such as blurring, refining of lines and reduction of noise) to an image in preparation for classification by a ANN algorithm. The ANN will then usually compare regions of skin against those it already knows through supervised training using back-propagation. The problem with this method of classification however is that they rely on images of predetermined similar features, for instance a ANN trained to find images of whole body nudes will not work against images of partially or clothed people which some child pornography is classified.

Highlighting this last point brings about the question of how different child pornography is from adult pornography. Images of clothed adults would not be considered pornography unlike images of clothed children and calibrating a ANN to detect both clothed and unclothed children is not as simple as classifying by skin found in an image. If one were to ignore this as well and focus just on images of nude child pornography, the methods of ANN used in classifying adult pornography are also not adequate for classifying child pornography. The problem is that while established methods of classifying these images would work, they would not attempt to distinguish the differences between features of a child and those of an adult.

Alternatively, there has been research into classification by features within skin regions (Forsyth and Fleck 1999). This method is similar to the previous mentioned method where image data is processed in preparation for an ANN algorithm. The difference is what the ANN algorithm is trained not for skin regions within the image, but is classified by predefined lines or shapes within the skin regions. The intention with this is to classify an image by known nude body curvature or shapes that are invariant to size, pose, or translation.

Some researchers go on to say that colour alone is enough to identify pornography instead of going further with any specific pattern recognition such as body features. This may be true in the case of pornography alone, however when identifying people by age -colour alone does not provide enough information to make such a conclusion.

The problem is that in order to identify child pornography one needs more information than simply the quantity or shape of skin regions in an image. Ways of identifying human beings and their respective age from features is what is required.

While many methods of object recognition utilise either motion capture, sequential images, specialist equipment, or human interaction for human detection (Taylor 2000; Mori and Malik 2002; Liang, Tieniu et al. 2003; Remondino and Roidakis 2003; Hee-Deok, Park et al. 2007), other methods of object recognition explore some aspects of feature detection using only single image frames. One area of focus in this field is the use of scale invariant feature transform (SIFT) (Lowe 1987; Lowe 1991; Almeida, Rocha et al. 2008). SIFT detection seeks to locate feature vectors which over image scaling, rotation, translation, and partial illumination and geometric distortion changes are invariant to change. Key vector points are then compared against stored known object feature vectors to find a match.

Lowes, method (SIFT) is fast and robust, however is too precise to unique features and more useful for finding unique objects or people, than it is for determination of features such as age or height.

Age Identification

The field of age classification is one which some say falls under the area of ‘soft biometrics’ (Jain, Dass et al. 2004). Soft biometrics is a area of research dissimilar to traditional biometrics in that its goal is the identification of physiological features such as age, gender, hair skin and eye colour, height and weight, without needing close proximity to the target (hence the term ‘soft’ meaning loose). However, while many other features of the human physiology are feasibly calculable, it’s been claimed that no reliable methods exist for age verification in soft biometrics (Woodward 2000; Jain, Dass et al. 2004).

Automated age classification was first documented by researchers Kwon and Lobo (Young Ho and da Vitoria Lobo 1994). By reviewing studies in craniotomy, art, theoretical makeup, plastic surgery, and perception, Kwon and Lobo believed that the growth of the human head was the best source of age determination. They explained that through cardioidal strain transformation, the human head grows in a “series of ever growing circles all attached at a common tangent “base” point” at the top of the head, and followed that by analysing the growth of the lower parts of the face one could categorise a face into distinct age groups. By judging the ratios of distances between facial features, Kwon and Lobo’s produced an infant identification rate of just under 68% with their algorithm. Additionally Kwon and Lobo’s sample size was limited to 47 images. However, following Kwon and Lobo’s methods researchers have improved the success rate to over 80% with over 4 times the sample size (Horng, Lee et al. 2001).

Possible solution and problems

Reviewing the literature established so far one can conclude that a combination of age classification by facial analysis, and an application of flesh detection such as Forsyth and Fleck solution within close proximity to identified faces of children, one may be able to produce a method of classifying child pornography. Loosely a solution would operate in the manner illustrated in Figure 1:

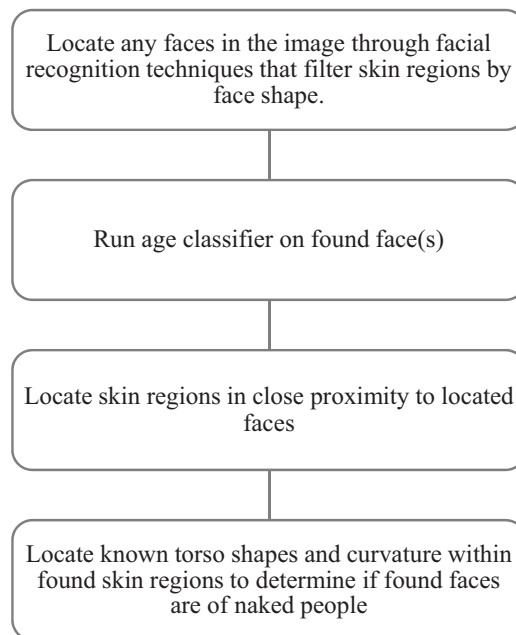


Figure 1: Child pornography detection based on existing research. Method attempts to apply theories based on age verification by facial features and later attempts to assess naked body parts using recognition of shapes and lines similar to those located on the body, namely within the torso region.

The question would remain whether the method would hold low false positive and high success rate until tested, however another question that can be made is how robust would such a method be.

Presentable issues with this solutions' results may be the elements that impede the continuation of the process, such as amount of content available in the scene. For example as the face classifier may confirm a face of a young person there may not be enough information present to confirm the degree of nakedness such as interference by clothing, objects, or image clipping. Additionally, factors such as camera quality, camera viewpoint, environmental lighting, and image noise or compression artifacts can also impede the processes of locating faces, body features, or object segmentation and grouping.

However the greatest issue making such a solution flawed is that it does not actually seek to classify imagery as child pornography. As stated earlier in this paper, an image of child pornography can be varied in depiction. One image of child pornography may display a clothed child in a sexual pose or being assaulted, while another may represent a child completely naked being sexually abused. However, an image of a children and family naked being represented in a form of naturism would not be classified as child pornography.

RESEARCH METHODOLOGY

Research Questions

The hypothesis of this research is "Child pornography can be automatically identified and classified from single-frame images by a computer system".

In order to prove this hypothesis the following research questions are posed:

- What aspects of computer vision theory are most important in the process of identifying child pornography?
- How robust can a automated child pornography detection system be made?

Research Design

The proposed research is set to be split into

- Creation and testing of a person recognition and segmentation system.
- Creation and testing of a pose and body explicitness system.
- Creation and testing of a age verification system.
- Testing of combined system with actual child pornography.

Phase 1 – Creation and testing of a person recognition and segmentation system.

After developing a method of identifying and segmenting whole people from images, clothed or otherwise, the system will need to be tested against a large sample of images for both robustness against factors of image resolution and detail, noise and compression artifacts, and scene interference or clothing.

Phase 2 – Creation and testing of a pose and body explicitness system.

Once phase 1 is able to reliably segment whole people in a scene, a separate system needs to take the segment or person, measure pose, and check for explicitness such as undergarments, flesh, or visible extremities and genitalia. Once a system is established the same testing methodology must be applied from phase 1 for robustness, however the image selection will use normal adult pornography ignoring any facets of age verification and focusing on classification of pose and nudity.

Phase 3 – Creation and testing of a age verification system.

Similar to phase 2, phase 3 relies on the use of phase 1's segmentation. With this a system must be created to locate the face and or height of the subject and make a classification of age. Height may be excluded unless other human subjects are located in the scene to compare against. Once the system is established the same testing methodology from phase 1 is

applied to test robustness of the system. Image selection will use legitimate images of children instead of adults to classify age.

Phase 4 – Testing of combined system with actual child pornography.

Once phase 2 and 3 are successful, both systems are merged and are blind tested against a database of actual child pornography. Reliance of this stage would require cooperation with law enforcement to perform the testing on behalf of the researcher to adhere to the code of federal law while being able to find true results of the system.

CONCLUSION AND FURTHER RESEARCH

Having established that existing research into tradition pornography and age verification is not enough to confirm the existence of child pornography, another solution is required.

Such a solution may adapt established theories of computer vision to analyse the image as a whole. The system would need the ability to ignore the constraints of an uncalibrated image, retrieving as much information as possible and making an overall determination of object segmentation, and grouping to be able to determine the existence of people, pose, explicitness, and most importantly age.

As research continues, more research needs to be known about general computer vision. It's hypothesised by this researcher that no one theory or method will be significant enough to classify child pornography but a combination of computer vision techniques may develop one significant enough to eventually assist law enforcement locate evidence, prosecute criminals, and save lives.

REFERENCES

- Almeida, J., A. Rocha, et al. (2008). Making colors worth more than a thousand words. Proceedings of the 2008 ACM symposium on Applied computing. Fortaleza, Ceara, Brazil, ACM.
- Biegel, S. (2001). Beyond our control? Confronting the limits of our legal system in the age of cyberspace. Cambridge, MA, MIT Press.
- Forsyth, D. A. and M. M. Fleck (1999). "Automatic Detection of Human Nudes." International Journal of Computer Vision(Volume 32, Number 1 / August, 1999): 63-77.
- Hee-Deok, Y., A. Y. Park, et al. (2007). "Gesture Spotting and Recognition for Human–Robot Interaction." Robotics, IEEE Transactions on **23**(2): 256-270.
- Hong, W., C. Lee, et al. (2001). "Classification of age groups based on facial featuresq." Tamkang Journal of Science and Engineering.
- Jain, A. K., S. C. Dass, et al. (2004). Can soft biometric traits assist user recognition? SPIE.
- Jain, A. K., S. C. Dass, et al. (2004). Soft Biometric Traits for Personal Recognition Systems.
- Jenkins, P. (2001). Beyond tolerance: Child pornography on the Internet. New York, New York University Press.
- Jones, M. J. and J. M. Rehg (2002). "Statistical Color Models with Application to Skin Detection." International Journal of Computer Vision(Volume 46, Number 1): 81-96.
- Liang, W., T. Tieniu, et al. (2003). "Silhouette analysis-based gait recognition for human identification." Pattern Analysis and Machine Intelligence, IEEE Transactions on **25**(12): 1505-1518.
- Lowe, D. G. (1987). "Three-Dimensional Object Recognition from Single Two-Dimensional Images." Artificial Intelligence **31**(3): 355–395.
- Lowe, D. G. (1991). "Fitting parameterized three-dimensional models to images." IEEE Transactions on Pattern Analysis and Machine Intelligence **13**(5): 441-450.
- Marshall, W. and S. Barrett (1990). Criminal neglect: Why sex offenders go free, Doubleday Canada.

Mori, G. and J. Malik (2002). Estimating Human Body Configurations Using Shape Context Matching Computer Vision — ECCV 2002: 666-680.

Remondino, F. and A. Roditakis (2003). Human figure reconstruction and modeling from single image or monocular video sequence. 3-D Digital Imaging and Modeling, 2003. 3DIM 2003. Proceedings. Fourth International Conference on.

Taylor, C. J. (2000). "Reconstruction of articulated objects from point correspondences in a single uncalibrated image." Comput. Vis. Image Underst. **80**(3): 349-363.

Taylor, M. and E. Quayle (2003). Child Pornography: An Internet Crime, Psychology Press.

Wells, M., D. Finkelhor, et al. (2007). "Defining Child Pornography: Law Enforcement Dilemmas in Investigations of Internet Child Pornography Possession." Police Practice and Research: An International Journal **8**(3): 269-282.

Woodward, J. D. (2000). "Testimony to the Commission on Online Child Protection on Age Verification Technologies." from <http://www.copacommission.org/meetings/hearing1/woodward.test.pdf>.

Young Ho, K. and N. da Vitoria Lobo (1994). Age classification from facial images. Computer Vision and Pattern Recognition, 1994. Proceedings CVPR '94., 1994 IEEE Computer Society Conference on.

Zheng, Q.-F. and W. Gao (2005). Fast Adaptive Skin Detection in JPEG Images. Lecture Notes in Computer Science. Heidelberg, Springer Berlin. **3768**: 595-605.

Zhu, Q., C.-T. Wu, et al. (2004). An adaptive skin model and its application to objectionable image filtering. Proceedings of the 12th annual ACM international conference on Multimedia. New York, NY, USA, ACM.

COPYRIGHT

Glen Thompson ©2009. The author/s assign the Security Research Centre (SECAU) & Edith Cowan University a non-exclusive license to use this document for personal use provided that the article is used in full and this copyright statement is reproduced. The authors also grant a non-exclusive license to SECAU & ECU to publish this document in full in the Conference Proceedings. Such documents may be published on the World Wide Web, CD-ROM, in printed form, and on mirror sites on the World Wide Web. Any other usage is prohibited without the express permission of the authors.