Facial Recognition in Low-Light

Kwesi Byrd1, Hiba Al-kawaz1, Nathan Clarke1,2 and Steven Furnell1,2,3 1School of Computing, Electronics and Mathematics, University of Plymouth, United Kingdom 2Security Research Institute, Edith Cowan University, Perth, Western Australia 3Centre for Research in Information and Cyber Security, Nelson Mandela University, Port Elizabeth, South Africa info@cscan.org

Abstract

It has long been the goal of facial recognition systems to be able to operate with high success rates in any environment regardless of illumination. Many researchers have tried to overcome the issue of illumination both algorithmically and through the use of novel capture devices and techniques. Most of these novel techniques have some issue that usually prevent them from being successfully implemented in the real world, for example, the cost of implementation. This research project aims to examine to what extent, if any, can current commercially available cameras in the low to medium price range, overcome the issue of illumination variation in facial recognition. An experimental study was conducted utilizing the facial images of 30 live subjects, captured using four different devices across a range of low light settings. The images were then processed through three different facial recognition algorithms. The results indicate that better quality devices did outperform their cheaper counterparts across a range of variable lighting conditions, providing a significant improvement in the usability of the approach in practice.

Keywords: Facial Recognition, Low illumination, illumination, Biometrics.

1. Introduction

There are many reasons that facial recognition systems are one of the most frequently used biometric systems today. One reason is due to the fact that faces are very unique, highly complex, and extremely difficult to forge or recreate in any manner that would be acceptable for security purposes. In addition, facial recognition is one of the least intrusive methods of biometric identification. This results in facial recognition's high results for acceptability and collectability (Jain et al., 2004). However, facial recognition is not without its challenges and limitations with one of the main limitations being the issue of illumination. It has been well documented in literature that variations in illumination have a detrimental effect on the performance of facial recognition system (Zhao et al., 2003; Georghiades et al., 2001; Tan and Triggs, 2010; Kawade, 2002). This is particularly true with regards to the effect of low or dim lighting on the performance of various algorithms (Kawade, 2002). This began the search for a solution that could make facial recognition systems illumination invariant or "lighting independent". Many researchers have come up with possible solutions to overcome this issue, like infrared capture devices or 3D models of the face (Socolinsky and Selinger, 2002; Paysan et al., 2009). However, none of the novel capture devices or techniques are ready for implementation today or have a minimal cost of implementation. In order to meet the requirement of improving performance while simultaneously being ready for implementation and relatively low in cost, we should examine the current lexicon of ubiquitous, mid-priced digital cameras. A lot of digital cameras today purport to be able to deliver better image quality images in low light. Therefore, this research project sought to investigate to what degree currently available, moderately priced visible spectrum cameras are able to overcome the issue of poor illumination. Moreover, determine to what extent those better-quality images translate into better facial recognition performance.

The remainder of this research paper will be structured as follows: Section 2 provides an overview of existing research that has attempted to overcome the problem of illumination. The research methodology and results are presented in sections 3 and 4 respectively. A discussion of the results is presented in section 5, with the paper concluding with conclusion and future areas for further research.

2. Overview of illumination invariant facial recognition

One way that researchers have attempted to overcome the problem of illumination is by the use of images in the thermal infrared spectrum rather than the visible light spectrum. Since images captured in the thermal spectrum are based on internal temperature rather than reflectance of the light of the environment, images in the thermal spectrum are lighting independent (Socolinsky and Selinger, 2002). There are various novel techniques for facial recognition using thermal infrared images. One such technique proposed by in Buddharaju et al. (2007) involves using a thermal infrared camera for image capture and then uses an algorithm to extract a map of the superficial facial veins of the subjects face. Another technique created by Li et al.(2006), uses a combination of near infrared (NIR) LED's and a long pass optical filter to provide near infrared "illumination" filtering out all other light. Unfortunately, to date, these infrared technology solutions have not become cost efficient or widely available. The cost of infrared capture in facial recognition systems, compared to visible spectrum camera recognition systems, is what discourages real world implementation of the aforementioned techniques.

Besides infrared capture devices, researchers have also attempted to overcome the problem through the use of novel algorithms and 3D modelling techniques. An example of one such algorithm is the Discrete Cosign Transformation or DCT algorithm. A DCT is a mathematical calculation that is performed on the pixel values of an image to alter them in a specific way (Ahmed et al., 1974). In regards to facial recognition, Chen et al. (2006) examined how a DCT could be used to "correct" the pixels of an image so that the image appeared to have been captured under ideal lighting conditions. Paysan et al (2009) explored the use of 3D models of the face to help provide better recognition performance. The aim of Paysan et al. (2009) was to create a 3D model that would save facial recognition researchers from having to create their own. This is because while 3D modelling is a powerful too, the models themselves are expensive and time consuming to produce. Both the DCT (Chen et al., 2006) and the 3D model created by Paysan et al. (2009) are further examples of the many techniques that have been created by researchers to overcome illumination issues.

While use of infrared and facial mapping have both recently made their debut as features within Apple's Face iPhone X and Face ID system (Apple, 2017), there are still many facial recognition implementations that depend upon more traditional camera technologies. As such, this research project (conducted just prior to the announcement of the iPhone X) aimed to investigate the degree that currently available, moderately priced visible spectrum cameras can help to overcome the issue of dim illumination, and to what extent the better quality images translate into better facial recognition performance.

3. Methodology

The design of this experiment was based loosely on an experiment by Kukula and Elliot (2004) with modifications to fit the scope of the aims of the study. The equipment setup and capture dimensions were similar to that of a mugshot or a passport photo. Participants were asked to look directly into the camera and have a neutral expression on their faces. The capture environment was a room completely devoid of natural light in order to tightly control the illumination in the environment. Two separate capture environments were prepared in this manner for the convenience of the study participants.

Five different light configurations were selected in hopes of mimicking the enrolment and subsequent verification of individuals in an environment with less than ideal lighting conditions. In the first 3 settings, a lamp was used to light the face in an ambient manner at progressively lower light settings starting at around 7.0 lux, then around 5.0 lux and lastly around 3.5 lux. For the remaining 2 light settings, the same lamp was used to directly light the face at progressively lower light settings also starting at about 7.0 lux and ending at approximately 5.0 Lux. The light produced by the lamps was measured by a digital lux meter.

The lux meter was used to measure the light at the face of the subject before each group of images were captured. The lux meter was also used to measure the light in the capture environment prior to any illumination adjustments being made or images being captured to ensure the capture environment was devoid of any other light (natural or artificial) that would affect the measurement of data.

Four different capture devices were used to capture all of the images in the experiment. The chosen devices were all part of the lexicon of ubiquitous, commercially available and relatively cheap visual spectrum capture devices. The price range of the cameras used was between 70 and 275 pounds. The average megapixel range was between 8 and 20 megapixels. An image was captured by each of the 4 cameras at each of the 5 light settings for a total of 20 images per person. As there were 30 participants, this made for a total of 600 images. All participants were non-vulnerable adults over the age of 18. All subjects were of white European origin. Gender distribution of the participants was 57% female and 43% male.

The images were tested using systems based on 3 different algorithms. The first two systems was based on the Fisherfaces and Fourier-Bessel algorithm. The third algorithm was Amazon's Rekognition service. Whilst are all commercial, the first two algorithms are Matlab implementations of research articles, rather than a fully-fledge commercial system such as Rekognition. The selection of algorithms were chosen to investigate the performance of differing biometric classifiers and to understand the performance of commercial/industry-grade software. The database of each system consisted of one training image per participant taken at the highest illumination setting. Prior to training and testing, the images were put through pre-processing using the Viola-Jones algorithm to normalize the images with regards to size, aspect ratio and orientation without the images appearing compressed or stretched.

4. Results

The first two observations in this experiment came in midst of data collection. The image captured by each camera appeared drastically different in terms of brightness and clarity when the lux measurement and subject remiained the same. Figure 1 below is an illustration of that. The 4 images below were captured in the same environment at a light setting of 5.1 lux.

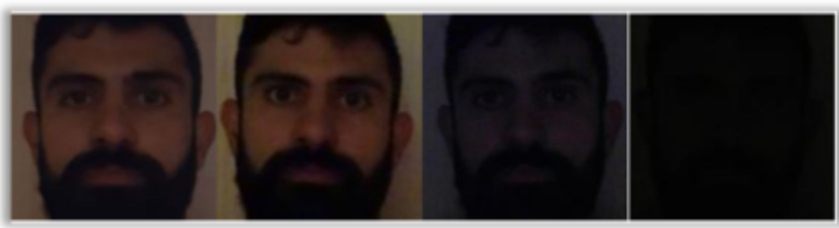

Figure 1: All images captured at the 5.1 lix

Secondly, the original experiment was designed to have a total of 5 different capture devices. However, it was discovered during the initial data capture that the Intel SDK camera could not capture photos at the lower light settings needed for this experiment. When an image capture was attempted with this camera, a black image resulted with no recognizable facial image. Further image capture was attempted in both of the different capture environments in order to rule out other environmental issues. Upon further investigation, it was discovered that the camera ceases to be able to capture images at anything lower than 14.9 lux. The results of the research are presented in Table 1 and Table 2.

The results are broken down into performance by light setting and performance by device. Table 1 shows, independent of classifier, a significant difference in performance across devices, with the Sony QX30 achieving a 79% recognition in comparison to the worse performer, the Logitech C920, with a performance of 3.3%. This would suggest the quality of the camera and the associated image processing engines which comprise them can have a significant impact on the recognition.

Table 1: Results by capture device

	Sony QX30	Galaxy S6	Microsoft Lifecam	Logitech C920
FisherFaces	37.3%	30.7%	6%	3.3%
Fourier Bessel	41.3%	46.7%	8.0%	3.3%
Amazon Rekognition	90%	83.20%	38.40%	5.60%

Table 2: Results by light setting

	Setting 1 (7.0 lux ambient)	Setting (5.0 lux ambient)	Setting 3 (3.5 lux ambient)	Setting 4 (6.5 lux direct)	Setting 5 (5.0 lux direct)
FisherFaces	37.5%	20%	10%	17.5%	12.5%
Fourier Bessel	44.2%	25.8%	13.3%	24.2%	16.7%
Amazon Rekognition	55.40%	63%	42%	57%	45%

An analysis of the performance of differing capture devices utilizing Amazon Rekognition (as this was the best performing algorithm amongst the classifiers) did show an interesting correlation between recognition performance and the quality of the camera – with the more expensive camera's performing significantly better than their cheaper counterparts. This strongly suggests that modern fully functional digital cameras do contain sufficient additional technology to enable successful capture across a wider range of light settings. It should be noted, in order to perform this analysis, images from the Sony QX30 camera, light setting 1 were used as the enrollment dataset. These were selected based upon the premise of ensuring the enrollment template is based upon the best quality sample; however, it also leads to an artificially high classification result for that particular set of verifications. Nonetheless, this weakness does not impact the general trend that can be seen in the quality of the recognition performance across light settings.

Table 3: Results by capture devices per each setting (Amazon Rekognition)

	Setting 1 (7.0 lux ambient)	Setting (5.0 lux ambient)	Setting 3 (3.5 lux ambient)	Setting 4 (6.5 lux direct)	Setting 5 (5.0 lux direct)
Sony QX30	100%	92%	84%	92%	92%
Galaxy S6	88%	88%	76%	84%	80%
Microsoft Lifecam	68%	68%	4%	44%	8%
Logitech C920	12%	4%	4%	8%	0%

5. Discussion

Based on examination of the results, it was concluded that modern, moderately priced cameras have made a significant improvement over their cheaper/older counterparts in being able to overcome the issue of low illumination. With all of the results across light settings and devices, a pattern emerged in the performance of the cameras. The Galaxy S6 and the Sony QX30 largely demonstrated a downward slope in performance as the light in the capture environment decreased. The Logitech and the Microsoft LifeCam, however, consistently performed at the bottom of the metrics across all light settings. This is most likely due to the Logitech and Microsoft web cameras not having as much control over ISO, aperture size and shutter speed as the other 2 capture devices.

The general recognition performance on the FisherFaces and Fourier-Bessel approaches were very poor. This was largely due to the quality of the underlying technique. Amazon's Rekognition approach provided a far more industry-comparable result. Nevertheless, all classifiers demonstrated the same trend in reducing performance across light settings – indicating that classifiers are beginning to struggle (even when the capture device is capable of advanced image processing). This offers up the opportunity for further work to be undertaken to explore the nature of the feature vector to identify whether particular features are more easily observable in low light and if so, what value they have in terms of discriminative ability. This would allow the development of facial recognition algorithms that are designed to utilize feature vectors that can be extracted from low-light captures, rather than being dependent upon a complete feature vector whose composition could be significantly challenging in low light.

6. Conclusion

In conclusion, a possible solution to the issue of illumination variance in facial recognition was put forth and tested. An experiment was conducted using currently available visible spectrum cameras to ascertain whether or not cameras with improved lowlight capture technology could solve the problem of illumination variance in facial recognition. 600 images were captured from 30 subjects. The results produced proved that current mid-priced commercially available cameras are able to solve the issue of illumination to a good degree whilst still using current facial recognition algorithms.

The latest devices, such as the iPhone X, have shown that low-light limitations can be overcome by integrating infrared camera and illumination technologies. However, this comes at a cost of integrating additional technology that may not have any other purpose within the device. In cases where it is desired to leverage an existing., traditional camera, low-light considerations still remain. In addition, further areas for future research will seek to explore the remaining issues and barriers that exist within facial recognition including facial pose, the use of accessories and orientation of camera and subject. Overcoming these issues will provide a more usable and transparent approach that can be applied in a number of contexts, from implicit authentication on mobile devices, to better suspect identification from CCTV in digital forensics.

References

Ahmed, N., Natarajan, T., & Rao, K. R. (1974). Discrete cosine transform. IEEE transactions on Computers, 100(1), pp.90-93.

Apple. 2017. Face ID Security. November 2017. https://images.apple.com/business/docs/FaceID_Security_Guide.pdf

Buddharaju, P., Pavlidis, I. T., Tsiamyrtzis, P., & Bazakos, M. (2007). Physiology-based face recognition in the thermal infrared spectrum. IEEE transactions on pattern analysis and machine intelligence, 29(4).

Chen, W., Er, M. J., & Wu, S. (2006). Illumination compensation and normalization for robust face recognition using discrete cosine transform in logarithm domain. IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), 36(2), pp.458-466.

Georghiades, A. S., Belhumeur, P. N., & Kriegman, D. J. (2001). From few to many: Illumination cone models for face recognition under variable lighting and pose. IEEE transactions on pattern analysis and machine intelligence, 23(6), pp.643-660.

Jain, A. K., Ross, A., & Prabhakar, S. (2004). An introduction to biometric recognition. IEEE Transactions on circuits and systems for video technology, 14(1), pp.4-20.

Kawade, M. (2002). Vision-Based Face Understanding Technologies and Applications. In: Micromechatronics and Human Science, 2002. MHS 2002. Proceedings of 2002 International Symposium on. [online] IEEE. Available at: http://ieeexplore.ieee.org/document/1058006/ [Accessed 1 Mar. 2017].

Kukula, E. P., & Elliott, S. J. (2004). Evaluation of a facial recognition algorithm across three illumination conditions. IEEE Aerospace and Electronic Systems Magazine, 19(9), pp.19-23

Li, S. Z., Zhang, L., Liao, S., Zhu, X., Chu, R., Ao, M., & He, R. (2006). April. A near-infrared image based face recognition system. In Automatic Face and Gesture Recognition, 2006. FGR 2006. 7th International Conference on (pp. 455-460). IEEE.

Paysan, P., Knothe, R., Amberg, B., Romdhani, S., & Vetter, T., 2009, September. A 3D face model for pose and illumination invariant face recognition. In Advanced Video and Signal Based Surveillance, 2009. AVSS'09. Sixth IEEE International Conference on (pp. 296-301). IEEE

Socolinsky, D., & Selinger, A. (2002). A comparative analysis of face recognition performance with visible and thermal infrared imagery. In: Pattern Recognition, 2002. Proceedings. 16th International Conference on. [online] IEEE. Available at: http://ieeexplore.ieee.org/document/1047436/ [Accessed 15 Feb. 2017].

Tan, X., & Triggs, B. (2010). Enhanced local texture feature sets for face recognition under difficult lighting conditions. IEEE transactions on image processing, 19(6), pp.1635-1650.

Zhao, W., Chellappa, R., Phillips, P.J., & Rosenfeld, A. (2003). Face recognition: A literature survey. ACM computing surveys (CSUR), 35(4), pp.399-458.