A Literature Review on the Technical Aspects of Facial Recognition Technology and Its Societal Impact

A Technical Report submitted to the Department of Computer Science

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Thinh Tu
Spring, 2021

Technical Project Team Members
Thinh Tu

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Signature _____Thinh Tu_____ Date __5/12/21___
Thinh Tu

Approved ____Hongning Wang_____ Date __5/6/21__
Hongning Wang, Department of Computer Science

A Literature Review on the Technical Aspects of Facial Recognition Technology and Its Societal Impact

Thinh Tu
Department of Computer Science
University of Virginia
tct2tf@virginia.edu

ABSTRACT

Facial recognition technology's capabilities have expanded greatly in recent years, bringing with it many different forms of application, whether it be for education, health, commercial uses, or security. Development and integration of facial recognition into different aspects of daily life gives the technology and system increasing agency and influence. This study seeks to explore the development of facial recognition technology through an analysis of research in the field. This literature review would focus on the combined contributions of different research done in the field of facial recognition, and how those contributions have helped moved the field forward or paved the way for new research. A key goal of researchers in this field, which involves computer vision and various algorithms, is to create automated face recognition systems that can equal, and eventually surpass, human performance. In order to achieve this, knowledge of the key findings from experimental studies of face recognition would be crucial. Another important issue to consider is the fact that development and integration of facial recognition into different aspects of daily life gives the technology increasing agency and influence. However, the growth of facial recognition is not without conflict and controversy, depending on the society and culture in which the technology is developed. As such, this literature review will also explore social aspects of the technology as well.

INTRODUCTION

Facial recognition systems are built on computer programs that analyze images of human faces for the purpose of identifying them. Facial recognition technology is developing at a rapid pace in the past decade, allowing for many novel applications that can improve general quality of life, provide new and more advanced medical treatments, facilitate the commercial interactions between the provider and its customer, as well as enhance law enforcement capabilities and improve public safety, amongst its many other possible uses, which are growing everyday as the technology continues to advance. The general goal of researchers and developers in the field of facial recognition is to continuously improve upon existing technologies in order to surpass human performance in terms of accuracy as well as speed and efficiency. With this in mind, this literature review seeks to aggregate and analyze the combined contributions of various research done in this field in order to facilitate further exploration in the development of facial recognition systems. However, it is also important to note that the growth and propagation of the technology all around the world is not without various ethical and social concerns, especially in regards to the way that the technology's capabilities could be applied for unethical purposes or without the consent of individuals subjected to the technology's usage. Even in cases of ethical application of facial recognition technology, the societal implications and impact can be very significant, to the point of changing established social and cultural norms. This literature review will also seek to explore some societal impacts as well as ethical concerns regarding the continued improvements and spread of the technology.

BACKGROUND

Although facial recognition technology may be thought of as fairly recent in the field of computer science, automatic facial recognition's beginnings dates back to the 1960s, when computers were first used in conjunction with humans input in order to distinguish human faces based on facial features. The technology received financial support from various governmental organizations to assist security, intelligence, and law enforcement personnel in the performance of their duties. Research on facial recognition to reliably locate a face in an image that contains other objects, as opposed to images only containing faces being used in earlier versions of the technology, gained traction in the early 1990s with the application of principle component analysis using eigenfaces. This was further improved upon using linear discriminant analysis, a generalization of Fisher's linear discriminant, a method used in statistics and other fields, to find a linear combination of features that characterizes or separates two or more classes of objects or events, which can be used as a linear classifier. Facial recognition algorithms are generally divided into two main approaches: geometric, which looks at distinguishing features, or photo-metric, which is a statistical approach that distills an image into values and compares the values with templates to eliminate variances. These algorithms can be classified into two broad categories: holistic and feature-based models. The holistic models attempt to recognize the face in its entirety while the feature-based models divide an image into smaller components like notable features and analyze each feature as well as its spatial location with respect to other features.

RELATED WORK

As facial recognition is a broad and significant field of study, various literature reviews have been done with a similar purpose of facilitating the development or providing a better general understanding of the field. “Reliable Face Recognition Methods - System Design, Implementation and Evaluation” written by Harry Wechsler [2] sought to comprehensively address the face recognition problem while gaining new insights from related and complementary fields, including neurosciences, statistics, signal and image processing, computer vision, machine learning and data mining. The book examines the evolution of research surrounding the field to date, explores new directions, and offers the author’s thoughts on the most promising areas for future research and development. “Face Recognition: A Literature Survey” by W. Zhao et al. [7] provides an up-to-date critical survey of still- and video-based face recognition research, including a review of the existing literature, and insights into the studies of machine recognition of faces. This paper explores and categorizes existing recognition techniques and provides descriptions of representative methods within each category. The review also consists of complementary and relevant topics such as psychophysical studies, system evaluation, and issues of illumination and pose variation. While these literature reviews are very comprehensive in regards to their coverage of the various technological aspects of the field of facial recognition, in general, there is little focus on the societal aspects of the technology’s development and usage. As such, this literature review will seek to bridge the gap between the aforementioned topics.

LITERATURE REVIEW

This literature review will consist of a general overview of current facial recognition techniques and the technologies involved, as well as some advantages and disadvantages of each method. It will also consist of an exploration of the societal impacts and common ethical concerns raised surrounding the field of facial recognition. As other literature reviews already provide a fairly comprehensive description and valuable discussions regarding the various techniques as well as the evolution of research around the technical development of the field, more than can be viably done given the time and individual constraints of this project, this review will synthesize and summarize the findings of Tolba, Ahmad & El-Baz, Ali & El-Harby, Ahmed [6] in Face Recognition: A Literature Review, published in the International Journal of Signal Processing, and Face Recognition by Humans: Nineteen Results All Computer Vision Researchers Should Know About by Sinha et al. [5] published by the Institute of Electrical and Electronics Engineers, as a general overview of the field, before exploring societal and ethical implications and concerns through Andrejevic and Selwyn’s [1] discussion of critical questions that concern, including but not limited to issues that arise due to the application of facial recognition technology in schools. Ethics of facial recognition will be explored through an analysis of Martinez-Martin’s [4] discussion of the important ethical implications of using facial recognition technology published in the AMA Journal of Ethics.

1 Facial Recognition as Done by Humans

Extensive research has been done with the purpose of creating a facial recognition system that can be deployed in any setting, regardless of viewing distance, illumination, sensor noise, or any other factors that may negatively impact the accuracy of such a system. However, as of this point in time, while systems may be able to outperform their human counterparts in different situations and in areas such as efficiency, no system is considered to be absolutely superior to the human visual system.

Research done by Sinha et al. [5] found that humans can recognize familiar faces in low-resolution images. Progressive improvements in camera technology have enabled essentially any device to take high resolution photos, which allow recognition systems to discriminate between individuals based on fine differences in facial features. However, in situations where high resolution images are not available, such as in situations where individuals have to be recognized from a distance like in public CCTV footage, such detail-based systems may struggle in terms of accuracy. Humans, on the other hand, are able to distinguish more than half of an unprimed set of familiar faces that had been blurred to have equivalent image resolutions of just 7×10 pixels. This result demonstrates that fine details are not absolutely necessary to achieve good facial recognition performance. Furthermore, this study found that the ability to tolerate degradations in face images increases dramatically with the familiarity to the subject of the image. Human observers’ performance suffers when asked to match two different low- resolution photographs of an unfamiliar person, and recognition is much better when surveying video footage of familiar colleagues as opposed to those with whom the observers have interacted with infrequently. With these results, one may consider that the increased performance related to familiarity may be attributed to recognition of other details besides facial features, such as body structure or gait information; However, this was shown to be mostly untrue. Recognition performance changes only slightly after obscuring the gait or body, but is affected dramatically when the face is hidden, indicating that while gait or body structure may provide additional information in recognizing familiar subjects, it can’t be the sole reason for the dramatic difference in performance when observing familiar and unfamiliar individuals.

Another discovery that showcases the difference between performance of the human visual system and an automated computer vision algorithm is that high-frequency information may not be too crucial in effective facial recognition. In computer vision, edge maps are commonly used to capture what are thought to be the most important aspects of images, the discontinuities, while being invariant to differences in levels of illumination. Humans are also able to distinguish between and recognize faces from simple sketches are caricatures in a similar way, which seems to indicate the importance of spatial frequency information. However, further research in the field by Graham Davies and colleagues, showed that images which contain exclusively contour information are very difficult to recognize (specifically, they found that subjects could recognize only 47% of the line drawings compared to 90% of the original photographs. Bruce and colleagues expanded on this by showing that the depictions shown to the subjects also contained photometric cues, involving a face’s

photometric structure, which make human drawn faces much more recognizable compared to something like an edge map.

Studies considering the nature of how humans process facial features showed that human facial recognition is largely holistic. A study done by Sadr et al. [5] showed that after being shown just one feature, such as the eyes or, surprisingly, the eyebrows, can be sufficient for relatively accurate human facial recognition for many famous faces. However, when features on the top half of one face are combined with the bottom half of another face, the two distinct identities are very difficult to recognize. Interestingly, when the two halves of the face are misaligned, presumably disrupting normal holistic processing, the two identities are easily recognized. This seems to indicate that human feature processing is not piecemeal, but depends on the holistic context, which is markedly different from many feature-based computer vision models. This is further supported by the finding that for humans, eyebrows, which are important for conveying emotions and other nonverbal signals and can survive substantial image degradations, are one of the most important features for facial recognition along with the eyes. Another difference in the way humans and computer process facial features lies in the difference in the level of importance of the spatial relationship between each feature. Some computer vision systems involve precise measurements of attributes, including but not limited to, inter-eye distance, width of mouth, and length of nose [2], attributes which have been shown to not be particularly important for human facial recognition. Humans that were shown images where images are compressed or modified in such a way that the aforementioned attributes experience drastic changes had little to no loss in their ability to recognize the face. A set of attributes that remains unchanged after the compression are the ratios of distances within the same dimension, so computer vision systems may be able to utilize this method of distinguishing traits as well.

2 Facial Recognition Techniques

Facial recognition techniques that are widely employed and are considered in this literature review include eigenfaces, neural networks, dynamic link architecture, hidden Markov model, geometrical feature matching, and template matching.

2.1 Eigenfaces. As mentioned earlier, one of the early modern facial recognition techniques involved the use of principal component analysis for eigenfaces. Eigenfaces are the eigenvectors of the covariance matrix of the set of face images, which make up the principal components of the distribution of faces. This technique works on the argument that any face images could be approximately reconstructed by maintaining a collection of weights describing each face by projecting the face image onto the eigenpicture, which is a standard face picture. With this technique, each face can be represented exactly by a linear combination of eigenfaces, and can be approximated using eigenvectors with the largest eigenvalues. In one study considered, this method achieved 96, 85, and 64 percent accuracy over variations in lighting, orientation, and size, respectively [6]. In summary, eigenface appears as a fast, simple, and practical method. However, in general, it does not provide invariance over changes in scale and lighting conditions.

2.2 Neural Networks. Neural networks are also a popular technique for facial recognition, which might be attributed to its non-linearity in the network, which could lead to a more efficient feature extraction step compared to linear methods. Different constructions of neural networks are utilized for different aspects of facial recognition. A multilayer perceptron or convolutional neural network is generally used for face detection, while for face verification, a multi-resolution pyramid structure is used. The convolutional neural network used in one of the studies mentioned provides partial invariance to translation, rotation, scale, and deformation by extracting successively larger features in a hierarchical set of layers, and reported 96.2% correct recognition on ORL database of 400 images of 40 individuals. While recognizers based on neural networks are able to recognize up to 200 people and could achieve up to 96% correct recognition rate in approximately 1 second, the computing expense grows quickly when the number of classes (individuals) increases [6].

2.3 Graph Matching. Graph matching is another method of facial recognition that utilizes dynamic link architecture. Memorized objects are represented by sparse graphs, whose vertices are labeled with a multiresolution description in terms of a local power spectrum and whose edges are labeled with geometrical distance vectors. Elastic graph matching, which is performed by stochastic optimization of a matching cost function, is used for object recognition. One of the studies that utilized this method achieved recognition rates of 86.5% and 66.4% for the matching tests of 111 faces of 15-degree rotation and 110 faces of 30-degree rotation to a gallery of 112 neutral frontal views, showing that this technique is fairly effective in terms of rotation invariance [6]. An issue with this method is that the matching process is computationally expensive relative to other methods.

2.4 Hidden Markov Models. A Hidden Markov Model (HMM) is a statistical Markov model in which the system being modeled is assumed to be a Markov process with unobservable or "hidden" states. The HMM assumes that there is another process whose behavior depends on the Markov process. For the purposes of facial recognition, faces are divided into regions such as the eyes, nose, mouth, which can be associated with the states of an HMM. Each unknown test image is then matched against every HMM (each of which represents a different subject) and the match with the highest likelihood is considered the best match. A study using this approach achieved 87% accuracy using ORL database consisting of 400 images of 40 individuals [6]. However, this technique also suffers from expensive classification time and training time, and as such may not be viable for larger groups of individuals.

2.5 Geometrical Feature Matching. Geometrical feature matching techniques are based on the argument that overall geometrical configuration of the face features is sufficient for facial recognition. This is supported by results demonstrating that facial recognition is possible even at a resolution as low as 8 x 6 pixels, where facial features can't be seen in detail. The positions and size of the main facial features such as eyes, eyebrows, nose, mouth, etc. are described by a vector in geometrical feature matching. In one of the studies mentioned, 35 features were extracted to form a 35-dimensional vector, and recognition was performed using a Bayes classifier. This study achieved a recognition rate of 90% on a database of 47 people [6]. Other techniques using manually extracted distances were also able to achieve high matching accuracy as well.

2.6 Template Matching. Template matching is a technique in which a face from a single viewpoint is represented by either a single template, where the image is represented as a two-dimensional array of intensity values, or a set of multiple distinctive smaller templates representing important features. Test images are then compared to find parts of the image which match the template image(s). The cross-correlation output will be highest at places where the image structure matches the mask structure. In a study done by Bruneli and Poggio, template matching provided a 100% recognition rate on a database containing 188 images of 47 individuals [6]. This technique, while highly effective, suffers from high computational complexity, making it less viable for larger sets of images.

3 The Social and Ethical Implications of Facial Recognition Technology

While we have explored the differences between how the human visual system and common facial recognition systems process images and videos in order to achieve high accuracy, as well as possible areas of improvement, these sources have yet to consider social and ethical aspects in regards to the development of the technology. While research with the purpose of furthering the capabilities of the technology in terms of just its performance is certainly beneficial in terms of the growth of facial recognition, the technology has immense importance and agency in a socio-political context, and so consideration of these factors is meaningful as well.

One of the most widely used and argued for application of facial recognition is for security purposes. After the 9/11 terrorist attacks in the US, video surveillance and facial recognition systems became the subject of increased interest and controversy. In favor of face recognition technology, there is the promise of a powerful tool to aid national security. On the negative side, there are fears of an Orwellian invasion of privacy. Given the ongoing nature of the controversy, and the fact that facial recognition systems represent leading edge and rapidly changing technology, face recognition technology is currently a major issue in the area of social impact of technology. Ethical issues posed in the context of facial recognition for national security mainly revolve around violation of privacy, or the lack of detailed legal guidelines regarding implementation of the technology.

In the United States and the EU, backlash towards facial recognition comes in the form of accuracy concerns. A December 2019 National Institute of Standards and Technology (NIST) study evaluated the effects of factors such as race and sex on facial recognition software. The study analyzed 189 facial recognition algorithms from 99 developers, using collections of photographs with approximately 18 million images of eight million people pulled from databases provided by the US Department of State, the Department of Homeland Security and the Federal Bureau of Investigation. The study found disproportionately higher false positive rates for African American, Asian and Native American faces for one-to-one matching, and higher rates of false positives for African American females for one-to-many matching [3]. There is also concern about abuse of facial recognition by law enforcement agencies, leading to increased regulatory scrutiny. One example of abuse that was fairly significant was reports about Clearview AI, which counts many law enforcement agencies as clients, was found to have amassed more than three billion images scraped from publicly available social media websites. The company was allegedly collecting data without notice or consent.

Andrejevic and Selwyn examined ethical concerns regarding the burgeoning integration of facial recognition and facial detection into compulsory schooling to address issues such as campus security, automated registration and student emotion detection. They noted that so far, these technologies have largely been seen as routine additions to school systems with already extensive cultures of monitoring and surveillance. In countries such as the US, UK and Australia, these technologies have so far prompted little controversy or push-back [1]. They noted that the various benefits from the incorporation of the technology, such as more efficient and secure transactions, greater accountability, enhanced public safety and security, improved economic productivity, and commercial services may cause the public to openly welcome the technology. However, concerns being raised include issues of diminished accountability, compromised civil rights, and limitations on the concentration of power. In schools, Andrejevic and Selwyn argue that extensive incorporation of facial recognition could lead to ethical problems such as foregrounding of gender and race, dehumanization of schooling, increased authoritarian nature of schooling, and possible oppression of marginalized groups.

In terms of healthcare applications Martinez-Martin notes that there are also various different ethical dimensions of facial recognition technology. The use of facial recognition in health care suggests the importance of informed consent, data input and analysis quality, effective communication about incidental findings, and potential influence on patient-clinician relationships, so privacy and data protection are thought to present challenges for the use of the technology for health applications. There could also be ethical issues in terms of liability, if the software develops to the point where it can replace a physician's judgement [4].

CONCLUSIONS

In summary, there exist a wide variety of commonly used facial recognition techniques, each with their own advantages and drawbacks. Some of these techniques seek to imitate the way the human visual system performs facial recognition, while others are entirely unique to computers. What remains true in both of these cases is that there are still areas where the human visual system performs more effectively, and as such, ways of implementation with the human visual system in mind might be able to lead to the creation of a more efficient and effective computer vision technique that utilizes the strengths of both. However, in terms of the development of facial recognition technology moving forward, during the process of creating facial recognition technology techniques and algorithms, it might be prudent to also consider the social implications of the technology as well.

The various different concerns associated with each particular application of the technology illustrates the complexity of the ethics regarding facial recognition, and how there are different justifications for its extensive application in each case. The aforementioned cases are just some examples of areas in which facial recognition technology's growth may have a large impact, and the ethical considerations that follow. If possible, researchers

and developers could also consider the implementation of certain measures to prevent the abuse or unethical application of the technology as well, if there isn't sufficient regulation to prevent these harmful events from happening.

FUTURE WORK

This project was able to give a brief overview of the ways humans process visual cues and how this relates to human facial recognition, as well as an overview of commonly implemented techniques. From this, some possible areas of improvement for the technology were considered. However, due to time and manpower constraints, these areas might not be as well explored as they could be. As such, future research could delve into more key differences between humans and computer facial recognition, and also consider the viability of implementing certain human recognition techniques as algorithms. Ethical and social implications were also discussed and considered, but possible ways in which algorithms and techniques could be modified to possibly address some of these concerns would be good to consider in the future as well. This would include ways to prevent various problems that arise from current facial recognition models, such as issues with recognizing minorities or underrepresented populations.

REFERENCES

[1] Andrejevic, M. and Selwyn, N. (2019). Facial recognition technology in schools: critical questions and concerns. Learning, Media and Technology, 45(2), 115-128.

[2] Harry Wechsler. 2010. Reliable face recognition methods: system design, implementation and evaluation, New York: Springer.

[3] Jehl, L. and Prochaska, K. 2020. Public Backlash Calls Use Of Facial Recognition Systems Into Question. The National Law Review. Retrieved from https://www.natlawreview.com/article/public-backlash-calls-use-facial-recognition-systems-question

[4] Martinez-Martin N. 2019. What Are Important Ethical Implications of Using Facial Recognition Technology in Health Care?. AMA Journal of Ethics, 21(2), 180-187

[5] P. Sinha, B. Balas, Y. Ostrovsky and R. Russell. 2006. Face Recognition by Humans: Nineteen Results All Computer Vision Researchers Should Know About. Proceedings of the IEEE, vol. 94, no. 11, 1948-1962, doi: 10.1109/JPROC.2006.884093.

[6] Tolba, Ahmad & El-Baz, Ali & El-Harby, Ahmed. 2005. Face Recognition: A Literature Review. International Journal of Signal Processing. 2. 88-103.

[7] W. Zhao, R. Chellappa, P.J. Phillips, and A. Rosenfeld. 2003. Face recognition. ACM Computing Surveys 35, 4 (2003), 399–458. DOI:http://dx.doi.org/10.1145/954339.954342