

AI for Archives: Using Facial Recognition to Enhance Metadata

Florida International University

Rebecca Bakker, Digital Collections Librarian

Kelley Rowan, Digital Archives Librarian

Dr. Liting Hu, Assistant Professor

Boyuan Guan (Ph.D candidate), Application Developer

Pinchao Liu, Intern

Zhongzhou Li, Intern

Ruizhe He, Intern

Christine Monge, Intern

July 31, 2020

This project was made possible in part by a 2019 award from the [Catalyst Fund](#) at LYRASILS.



This research is licensed under a Creative Commons Attribution-ShareAlike 3.0 License
<https://creativecommons.org/licenses/by-sa/4.0/>

AI for Archives: Using Facial Recognition to Enhance Metadata

Background

Goal

The goal of this research project was to determine the most effective facial recognition applications that could be implemented into digital archive image collections from libraries, museums, and cultural heritage institutions. Computer scientists and librarians at Florida International University collaborated to conduct qualitative assessments of both face detection and face search using photographs from FIU's digital collections. Specifically, the facial recognition platforms OpenCV, Face++, and Amazon AWS were analyzed. This project seeks to assist LYRASILS community members who wish to incorporate facial recognition and other artificial intelligence technology into their digital collections and repositories as a method to reduce research time and enhance their collections with more complete metadata.

Introduction

This research seeks to address the long-standing challenges of incomplete metadata within archives and digital repositories, with the hopes of finding methods that can decrease the time involved in locating and matching images of varying archival subjects. This project builds upon previous research involving Artificial Intelligence (AI) methods in archival settings and aims to provide potential enhancements in regards to AI application within digital repositories.

When it comes to creating descriptive metadata for photographs, the process is often time-consuming and relies on the evanescent expertise of the curator. This knowledge is key when it comes to developing an insightful *and* accessible experience for the end-user. However, correctly identifying images of individuals within photographic archives is particularly labor-intensive and ultimately costly when it comes to the amount of time spent on processing a photographic collection; especially ones that have missing identifying metadata. Additionally, prominent individuals featured in collections may be known to archivists, librarians, and curators, but that resource is often lost when the institutional memory holders retire or leave. Lesser-known individuals within the repository may never be appropriately named if their identity is not quickly determined. Ultimately, in regards to preserving subject identity within repositories and archives, much of the information is heavily reliant on professional memory and achievable knowledge.

Preliminary Study

This study consisted of qualitative assessments of three facial recognition applications for accuracy and ease of use. By using photographs from digital collections held at FIU, the facial recognition software was analyzed in a real-world environment, including specific challenges

that come with digital archive images. This project sought to work primarily with facial recognition applications that were low-cost and easily accessible to librarians and archivists. Preliminary facial recognition applications that were reviewed included Adam Geitgey's Facial Recognition API, OpenFace, and Microsoft Azure Face. Ultimately, the project moved forward to focus on three facial recognition applications: OpenCV, Face++, Amazon AWS.

When it comes to the methodology of software such as OpenCV, Face++, and Amazon AWS, there are steps that the software utilizes to further analyze subjects. Face detection, which is the first step to facial recognition and analysis, refers to the ability of the software to identify the presence of a human face within a photograph or digital image. By using algorithms to search for facial features such as human eyes, nose, eyebrows, mouth, nostrils, and iris of the eyes, the software is able to present a thumbnail of what it assumes to be a human face. It can detect numerous faces within a group photo as well. Additional tests are then done to validate that the face that is separated is indeed human.

Once the software isolates a subject's face, facial recognition (a sort of identity recognition) ensues within the software. The process starts off with a human-curated database, in which human subjects are identified within photos so that the software can use that information to further identify other faces that are run through the program. Based on the database that is presented, the software can then analyze the face and essentially provide a similarity threshold with the images originally saved. Another step for recognition software is facial analysis, which is when the software itself analyses a subject's face for emotion, age, and gender. There are many differences in these software in regard to how many facial analyzation aspects there are; many software include factors such as age, gender, emotion, and things such as ethnicity and even skin quality. An additional feature is a confidence score, which lists the program's confidence level when it comes to finding a human face, or similarity within a database for a searched image.

When the research first began, database creation was incredibly important to allow adequate testing for the software being tested. The photographs used for this project came from the archives of the City of Miami Beach and the City of Coral Gables. All images used for this research were of photographs already uploaded and freely accessible to view in FIU's Digital Collections on [dPanther](#). Most were also aggregated to the broader platform of the Digital Public Library of America's repository. The images consisted of photographs of public figures or municipal officials, in their capacity as politicians, real estate developers, or city officials. The oldest photographs came from the 1920s, and the most recent from the 1990s. All the images were free of copyright restrictions.

For this project, 74 images were selected from two collections within FIU's digital repository as a software training set. The images included a mix of TIFFs and JPEGs. Thinking in advance to the facial recognition portion of the research, varying images of groups and individuals were featured, though most of them consisted of photographs of groups and crowds. A total of 248 faces were identified within the training set, most of which the subjects were unknown. Of the 248 faces, 40 of them were known by librarians. Those 40 were featured 145 times within the dataset and were manually cropped and labeled with their respective names. Of this group, 4 of them were women and 36 of them were men.

Software Tested

After much analysis, three different software were chosen for testing. Initially, we also tested Microsoft Azure, but the results were very similar to Amazon Rekognition, so we instead focused on software that had distinct benefits and drawbacks, and the results of Microsoft Azure are not included in this report. In the year since the research was proposed, there has been rapid development in the area of facial recognition software, and we believe these three applications to be the most appropriate and accessible programs for librarians and archivists to implement in a digital archive project.

OpenCV (Open Source Computer Vision Library)

OpenCV is an open-source software that can be used to facilitate various machine learning and computer vision actions. Serving as a platform for creation, it comes with a pre-established algorithm library that can easily be molded to an institution's preferences or work needs. One of the many pre-trained machine learning algorithms is a pre-trained deep learning face detector model, which can be trained and further edited through coding methods using Python, Java, C++, to name a few. It uses a local database rather than a cloud one like Amazon AWS and Face++. The deep learning face detection algorithm is the primary setup that will be used in the research for facial detection. This portion of the algorithm detects the human face, highlights it with a square, and lists the confidence percentage score right above it. For facial recognition, a local OpenCV model is used, working together with the facial detection aspect of it. Ideally, the facial detection model through OpenCV will scan varying facial features and structures, detect a face with a block around it, and list the confidence score as well as providing a face token within the code.

Amazon Web Services (AWS) - Rekognition

Amazon Rekognition is a cloud-based Software as a Service (SaaS) computer vision platform that was launched in 2016. AWS's facial recognition algorithm is able to identify elements such as mood, eyes open/closed, hair color, and other facial geometric features, and can create metadata tags for features such as a similarity threshold and can also create a Face and Image Identification, which is assigned by the software. The program does not focus on an individual's identity, but rather by the similarity of facial features. It can store 20 million different identities and has a maximum return search rate of 4096. AWS also uses its own security software by the name of Amazon Macie, which uses machine learning to enhance security around sensitive information (it also has its own separate cost). It has been sold and used by a number of United States government agencies, including ICE, Orlando and Florida police, as well as private entities. Using AWS in an archival and repository setting can prove to be helpful due to its metadata and ID creating elements. It also scans for things such as glasses, facial hair, emotion, and poses. AWS also offers a celebrity recognition aspect, which can detect celebrities through various settings—it can be helpful for identifying celebrities if any are found within a photo repository, though the extent of this celebrity database is unclear.

Face++

The most extensive of the three private software we tested was Face++, offering a wider range of detected information from subjects. Face++ is a cloud-based software that has various models used for human face and body recognition. This software provides data for elements such as ethnicity, skin color, and photo quality. Additionally, it boasts an unlimited amount of faces that can be stored within the program. These features seem to be the most beneficial when it comes to the vision of the program that is being created for the research - things like ethnicity, photo quality, and even skin color can help accentuate the similarity comparisons and can also prove helpful in the world of archiving, especially since a lot of the photographs being dealt with tend to be from times when photography wasn't so advanced. Face++'s extensive range of operations would provide useful for a repository or archive that is dealing with numerous photographic subjects that need identifying and grouping, because of its large facial storage. It also offers more recognition factors in contrast to Microsoft Azure and AWS. With more accessibility points when it comes to subject recognition, there is more opportunity to create a diverse facial data set with a controlled room for error.

Ethical/Privacy Concerns

When working with AI and facial recognition in particular, the ethical implications and privacy concerns that surround the technology had to be considered. This project remained focused on the use of AI within the use of digitized archives, and the study had a narrow focus and singular goal. However, we took several steps to ensure that - to the extent it was possible - the applications were used responsibly and ethically, as they would ideally be used in a repository setting. We were also mindful of recent trends regarding facial recognition software, and emerging best practices for using AI within libraries. We would encourage any practitioners interested in using these technologies within their own collections to review Thomas Padilla's *Responsible Operations: Data Science, Machine Learning, and AI in Libraries* and Ryan Cordell's *Machine Learning + Libraries A Report on the State of the Field* for further guidance.

The photographs used for this project came from the municipal archives, and as mentioned, were free of copyright restrictions. By choosing highly circulated and public images of officials, the research was able to use these subjects without having to worry that the person would expect privacy (due to their public personas and overall presence) and could therefore not object to their likeness being used for a project such as this one. However, because such care was taken to find appropriate images due to our lack of specific consent by the featured individuals, the dataset lacked diversity in regard to race, gender, and age. The majority of the photographs featured middle-aged to older white men. Though indicative of what local Miami politicians may have looked like at the time, this lack of diversity amplifies some of the chief concerns surrounding facial recognition software, and it's known issues with correctly identifying the faces of both women and people of color.

As we more recently know, many of these AI and facial recognition software are changing their standing about being deployed within law enforcement - Amazon AWS being one of them. Due to the recent protests and calls for defunding of the police by Black Lives Matter protesters and supporters, many AI software are pulling back their sales to law enforcement. As

stated by the *New York Times*, Amazon Recognition services are being pulled for a year from law enforcement applications, in order to give the United States Congress sufficient time to create regulation for ethical and lawful implementation of facial recognition software for civilian surveillance.

The use of facial recognition when it comes to criminal prosecution and monitoring can be tricky, and dangerous. As the *Washington Post* reports, “Privacy advocates have long raised concerns that police use of facial-recognition could lead to the wrongful arrests of people who bear only a resemblance to a video image. And studies have shown that facial-recognition systems misidentify people of color more often than white people.” These discrepancies in recognition can lead to wrongful convictions and arrests, and in most extreme cases uncalled police force on certain peoples. Implementation of AI facial recognition also begs the question of how ethically they are being used within police monitoring, and if they are being used as a way to spy on civilians and protesters.

For professional use, programs like Face++ and Amazon AWS, do save photographs to create a database for similarity checks when searching for an individual’s identity. There is a limit in regards to how many variations of a face and subject can be preserved. However, the question of if there are different variations of these softwares for policing still remains, and if these variations have a cap when it comes to storing faces and photographs. This also begs the question of if there are any limitations to these software (as there is when it comes to professional and private use), and how these software could be used in unethical privacy-infringement techniques by law enforcement. Despite the narrow scope of this research project, we found that ethical concerns of using AI, Machine Learning and Facial Recognition in the archived photographs had to be considered at almost every stage of our project.

Literature Review

Facial recognition has been recognized as a promising approach to efficiently identify needed materials with the least amount of staff time. Previous and continuing work on developing facial recognition suitable to archival quality photos are included in this literature review.

[Civil War Photo Sleuth](#) (CWPS) is a collaboration between the Virginia Center for Civil War Studies and the Crowd Intelligence Lab at Virginia Tech and Military Images Magazine and contains photos from the mid-1800’s on, making it a perfect example of using old, lower quality photos for facial recognition. It uses 27 facial points in order to identify soldiers despite changes in hair, hats, facial hair, and different photographic angles. As people upload photos, other mystery photos can be identified, making this a continually growing collection. In order to improve metadata, CWPS also makes use of crowd-sourcing techniques, such as tagging, and is linked to the Digital Public Library of America (DPLA) with hundreds of digital archives. Currently, they can identify a photo in mere seconds. Their goal is to be the world’s largest online archive of Civil War-era portraits, including soldiers, sailors and civilians.

[George Nott from Computerworld](#) reviewed efforts to develop and use facial recognition in archives, not only in the U.S. but in the Netherlands and Australia. Photo fit is a tool designed

to recognize and retrieve images of a person's face throughout the State Library of Australia's New South Wales' image collection and boasts a 95% confidence rate. Photo fit uses Amazon's Rekognition facial detection software and stores the results as JSON files in S3. However, as of 2019, further development was needed to achieve the results of attaching names to the images as found in Civil War Photo Sleuth.

In the Netherlands, Vintage Cloud, using the programming interface, Clarifai is doing similar work with machine learning models and facial recognition that improve accuracy. [Face-rec.org](https://face-rec.org) provides researchers with a list of nearly 100 databases dedicated to faces and facial recognition. Some are publicly accessible, while others, such as the MORPH Longitudinal Database, are for commercial and academic use. MORPH is the largest facial recognition database in the world containing 202,038 unique images of 40,395 subjects. Microsoft had the largest public facial recognition data set in the world (MS Celeb) with over 10 million images of approximately 100,000 individuals, but deleted it in 2019 as it likely ran afoul of the European General Data Protection Regulation laws.

Machine learning, as applied to facial recognition software, has made noteworthy advances during the current year. In their article, A Facial Expression Recognition Method Using Deep Convolutional Neural Networks Based on Edge Computing, authors Chen, Xing and Wang point out that of the three steps required for facial recognition, image preprocessing, feature extraction and facial expression classification, the last is the most problematic for current algorithms. To resolve this problem, the authors worked on developing deep neural networks, which adjust their behaviour and "learn" in ways similar to the human brain, making their computational power much faster and with better accuracy.

In July of 2020, scientists at the Graz University of Technology in Austria published an article addressing their research on [high energy consumption of artificial neural networks' learning activities](#). Researchers Wolfgang Maass and Robert Legenstein have developed an algorithm they call e-prop (e-propagation) that uses spikes in order to communicate between neurons in an artificial neural network. Additionally, working online rather than offline, helps manage energy usage during processing. This was a problem addressed in this grant by using a dedicated machine just for facial recognition processing. However, this method may make dedicated machines unnecessary and the research less expensive to conduct.

Methodology

In this research, we conducted face detection and face recognition with historical pictures from Florida International University's digital repository, dPanther. The research and testing of facial recognition software required developing both a photographic database and a cloud platform for the three facial recognition programs being studied: OpenCV, Face++, and Amazon AWS. The project's workflow and developing process can be viewed in Figure 1, and each step of the process will be illuminated below.

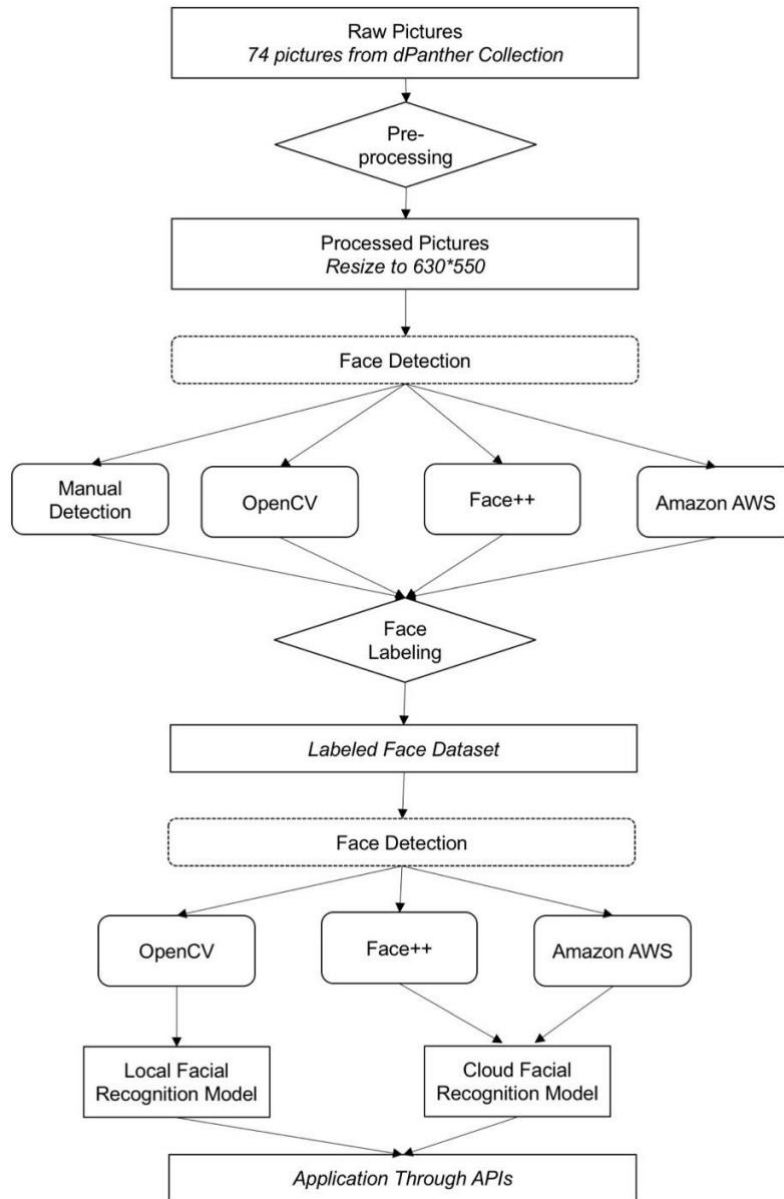


Fig. 1: Project Workflow and Development

Facial Detection

The first step of the project was to experiment with the facial detection function of each program. In order to build a baseline for this process, faces were first manually detected among the 74 raw photographs, as detailed in the Background section of this report. We then applied the face detection functions from OpenCV, Face++, and Amazon AWS respectively to the same dataset and compared accuracy results, as well as performance.

The second portion of testing was to build a facial recognition model by applying a face searching function. A face search allows the user to search the database for similar subjects. We used the three different facial recognition software to conduct this task, and compared their outcomes and performances. A labeled face dataset was also created during this task to be able to test the software's accuracy when it came to recognizing subject identities.

Ultimately, this step was divided into two branches: the local model build and the cloud model build. The local model build focused on OpenCV, as the program itself works with local databases. The cloud model portion pertained to Face++ and AWS, as these two software implement cloud-based libraries. For the local model, we tested with the pre-trained deep learning face detector model that comes within OpenCV's algorithm library. For the cloud model build, we uploaded our labeled face dataset to both the Face++ and AWS cloud platform and used their built-in models to conduct our experiments. A detailed look at the results of this portion of the experiment will be discussed below. The last task was to enable an application to access the facial recognition service built-in Task 2 by using RESTful APIs, or RESTful web service. It is our intention to develop a demonstration application in the dPanther platform that will implement this task in the future.

Initial Face Detection Experiment

In this experiment, we gathered 74 pictures from dPanther's archival photo collections of the City of Miami Beach and the City of Coral Gables. 248 faces from the 74 raw images were manually detected as the baseline for the experiment. The research then applied the face detection function from OpenCV, Face++, and Amazon AWS respectively to assess the performance of the three platforms:

- OpenCV software: Out of the 248 faces, it detected 170; however, only 152 of them were correct detections. It also must be noted that OpenCV had a significant time advantage. The total time it took to produce the results was 7.5 seconds for the entirety of the facial dataset.
- Face++ software: It had a superior performance advantage, with 183 of the 248 faces being detected. Within these 183 detections, there were no inaccurate detections. However, in contrast to OpenCV, it took longer to provide its results. Face++ would take 67.2 seconds to run through the dataset.
- AWS software: Upon its test, it returned 367 faces - well over the initial baseline that were manually detected. There were two detection errors: a missing detection

(Figure 2), and a false detection (Figure 3). Besides those two errors, the faces in the other pictures were appropriately recognized. The individual time spent per picture varied from 2 seconds to 7 seconds, and the overall time spent by AWS to provide results on the dataset was between 220 to 350 seconds.



Fig. 2: AWS's Missing Detection, showing faces of people (and background) not being detected

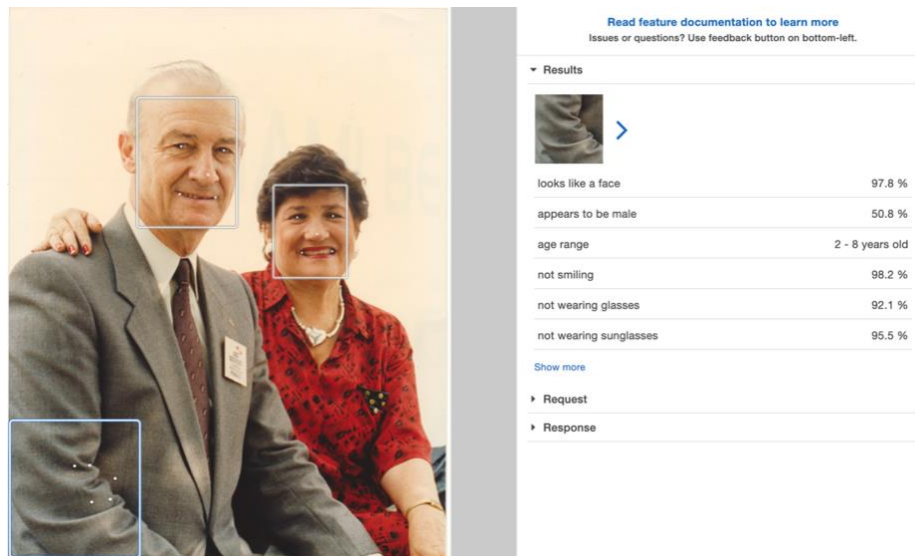


Fig. 3: AWS's False Detection, showing a detected face in an elbow

The final result of the Face Detection testing is summarized in Figure 4 below. In terms of faces detected, Amazon AWS outperformed the others with a total of 367 faces returned, surpassing the manual process of facial detection. It detected all the faces from background images such as paintings, artwork, and mirrors, which were not included in the manual process. Face++ came in next with the most faces detected, coupled with zero detection errors. Finally,

OpenCV was last with the least amount of faces detected, though it demonstrated a favorable amount of time spent in initial software facial detection, with the lowest processing time amongst the three.

	Face Detected	Time Spend (sec)	False Detection
Manual Detection	248	n/a	n/a
OpenCV	170	7.5	18
Face++	183	67.2	0
AWS	367	220-350	2

Fig. 4: Face Detection Summary Table

Face Recognition

After testing the face detecting capability of each software, tests were run on the three platforms to compare their face recognition and searching abilities. In total, there were 40 different people in the manually created training dataset of raw images. Many of the people in the actual photographs are either deceased, or, are otherwise difficult to find additional online images of them in traditional search engines, despite the fact they are/were public officials in Miami. Due to this, 28 pictures of nine different people who had a more predominant web presence were gathered via Google search. They are:

- Alex Daoud (4 pictures)
- Carl Fisher (4 pictures)
- David Dermer (6 pictures)
- George Merrick (4 pictures)
- Harold Rosen (1 picture)
- Matti Bower (3 pictures)
- Seymour Gelber (2 pictures)
- Simon Cruz (2 pictures)
- William Jennings Bryan (2 pictures)

All of the pictures were grouped into a separate folder by the person's name. Each individual face search across the programs would comb through the saved photos in the subjects' folder. This dataset would be used for the trails in Face++, OpenCV, and Amazon AWS. After each trail, the data output from the programs was written in a text file for comparisons.

The initial face recognition test was run on Face++. Using the created database, a face search was initiated. Face++ would return a confidence score based upon the similarity of the search and the found individual, as well as a face token (a unique identifier) per person. With the given face token, the research was able to verify whether it was an accurate outcome pertaining to the person in the original search.



Fig. 5: One search example using Alex Daoud: Most similar result (Left), searched online pictures (4 on the right)

```
Search 1:  
Original file:cruz_1619.jpg  
Name = Simon Cruz  
Confidence: 92.879  
Search 2:  
Original file:newpicDSC_0171_edited.jpg  
Name = Simon Cruz  
Confidence: 91.033  
Time used: 318ms
```

Fig. 6: Example output of a search

Face++ showed a great overall performance on accuracy and time used. Each individual search consisted from one to six photographs of the person. Out of the 28 searches, there were two errors, one of which was a false detection, and the other a missing detection. Other than these two errors, the remainder of the searches were correctly identified, with a confidence score for the accurate images averaging 84%. Each search consumed about 0.4 seconds. Since Face++ is based off of a Cloud API, there was no interruption within the computer's central processing unit. Utilizing the program only required a short period of internet access.

In testing OpenCV, there were a total of 8 instances with less than desirable results. Of the 28 photographs used, seven of them had a confidence value of less than 60%. Of those seven, OpenCV failed to recognize a subject. The average confidence of similarity suffered in comparison to Face++'s results, with OpenCV presenting a median confidence score of 68.1%. The speed of the facial searches within OpenCV averaged 3.1 seconds, which was a 2.7 second increase when compared to Face++. Considering that OpenCV is a local platform, prior to our testing it was estimated that the initial processing time would be less, however that was not the case.

Our testing of Amazon AWS presented the most favorable results. When inputting an image, AWS returned a confidence value based on the similarity of the two images being compared during the process. By default, when the confidence value is less than 90%, AWS will not consider the faces to belong to the same individual. Of the 28 images within the photographic dataset, AWS returned 24 correct results. Referring back to AWS's omission of those images that do not reach the threshold of 90%, three of them had confidence results of no less than 60%; therefore, not entirely disregarding the program's ability to accurately match images and individuals. The three had a confidence score of 82.06, 89.70, and 69.09. Ultimately, Amazon

AWS presented the highest average confidence percentage - a whopping 96.9%, topping both Face++ and OpenCV. However, the run time was an average of 31.8 seconds.

	Correct Comparison	False Comparison	Average Time spend(s)	Average Confidence
Face++	26	2	0.4	84
OpenCV	20	8	3.1	68.1
AWS	24	4	31.8	96.9

Fig. 7: Trail Outputs per Program

Lessons Learned

The two major approaches for facial recognition projects are an in-house trained model and cloud-based AI platform. For the experiments conducted in this study, we applied a sample photographic dataset from FIU's dPanther collection to both the in-house trained model (OpenCV) and cloud-base AI platform (AWS and Face++). Analyzing the outcomes from the study showed that in-house trained models need significantly more computing resources and more training data to obtain the same quality of results when compared to the cloud-based platforms. However, the in-house model outperformed the cloud-based platform in regards to time taken per transaction. Since OpenCV is locally housed, all actions were able to be performed without being reliant on network speeds. On the other hand, the cloud-based platforms such as AWS and Face++ gave a quick start and low-cost alternative for the average institutional archive looking to implement AI facial recognition software.

The tradeoff, however, is the scalability when a digital archives project is growing. Due to the fact that many archival image projects are always in rotation and aggregating new images, cloud-based API's may be an issue in regards to storage limitations and pricing. In addition to this, many archives and digital repositories deal with large images that can range in file size before initial editing. Services like AWS and Face++ cannot provide flexibility, especially for massive expansions within the database. They also rely heavily on the strength and speed of the network for its overall performance.

In general, for a small-scale project which does not have a large photographic dataset, or, in a situation where archivists and librarians do not have access to extensive IT resources, the cloud-based platforms are a good solution that provide an easy-to-start environment as well as an acceptable outcome for the project needs. However, for larger scale projects which have adequate IT resources as well as the continuously generated training data, an in-house trained model would be recommended. This will provide the most powerful and flexible environment for the project. Additionally, ethical issues regarding privacy are more pressing in the cloud-based

software, as the images are being used to train the data within a broader platform that provides less options for controlling the use of the images.

Future Development

Now that the testing of the facial recognition has concluded, the project is moving into an implementation phase. We are ready to experiment on how the software can be harnessed and used by digital archivists beyond the training set. In the coming months, we anticipate the creation of a sample application in the existing framework by applying Amazon Rekognition AWS's facial recognition capabilities to photographs in dPanther. This will better allow us to determine how the software works can be utilized by librarians and archivists, as they implement the selected facial recognition technologies into a digital repository workflow. This will not require any additional resources for the framework and will introduce the face detection and face search capability to the digital repository. We hope to highlight the development of the sample application in our upcoming webinar with LYRASIS. At the same time, we will continue to train the OpenCV model that we have implemented in our research. Once we have more training data, we will test the model and monitor any changes in accuracy.

References

- Civil War Photo Sleuth. (n.d.). Retrieved July 31, 2020, from <https://www.civilwarphotosleuth.com/>
- Castro, C. (2019). What's Wrong with Machine Bias. *Ergo, an Open Access Journal of Philosophy*, 6(20200523). doi:10.3998/ergo.12405314.0006.015
- Chen, A., Xing, H., Wang, F. (2020) A Facial Expression Recognition Method Using Deep Convolutional Neural Networks Based on Edge Computing. *IEEE Access*, 8. <https://www.doi.org/>
- Cordell, Ryan. (2020) Machine Learning + Libraries A Report on the State of the Field. Retrieved July 31, 2020 from <https://labs.loc.gov/static/labs/work/reports/Cordell-LOC-ML-report.pdf>
- Face Recognition Homepage. Retrieved July 31, 2020, from <https://www.face-rec.org/databases/>
- Graz University. (2020). New Learning Algorithm Should Significantly Expand the Possible Applications of AI. Retrieved July 31, 2020, from <https://techxplore.com/news/2020-07-algorithm-significantly-applications-ai.html?fbclid=IwAR20Rbbv4J1NUBa-11tWR2V3DfQDkpQyEJIAoP0GnOk0UTGUaPuSNqXq1B4>
- Greene, J. (2020, June 11). Microsoft won't sell police its facial-recognition technology, following similar moves by Amazon and IBM. *The Washington Post*. Retrieved July 31, 2020, from <https://www.washingtonpost.com/technology/2020/06/11/microsoft-facial-recognition/>
- Nott, G. (2019). Facial Recognition Tech Helping Place Faces From the Past. Retrieved July 31, 2020, from <https://www.computerworld.com/article/3487893/facial-recognition-tech-helping-place-faces-from-the-past.html>
- Padilla, Thomas. (2019) Responsible Operations: Data Science, Machine Learning, and AI in Libraries. Retrieved July 31, 2020, from <https://www.oclc.org/research/publications/2019/oclcresearch-responsible-operations-data-science-machine-learning-ai.html>
- Weise, K., & Singer, N. (2020). Amazon Pauses Police Use of Its Facial Recognition Software. *The New York Times*. Retrieved July 31, 2020, from <https://www.nytimes.com/2020/06/10/technology/amazon-facial-recognition-backlash.html#:~:text=SEATTLE%20%E2%80%94%20Amazon%20said%20on%20Wednesday,unfair%20treatment%20of%20African%2DAmericans.>