Exploring Inclusive Reasoning in AI

by

Charles V. Kittler, Olajumoke E. Jackson, Javier Marcos and Conrad C. Tulig

Worcester Polytechnic Institute

Advisor: Therese Smith

May 6, 2021

Abstract

Racism is embedded into the health industry in a subconscious level, influencing people in a subconscious level and backed by institutional policies. We set out to develop a machine learning algorithm to detect cancerous melanoma in a patient's skin, ensuring that the results of this system remain fair and consistent regardless of the patient's ethnicity or origin. While Convolutional Neural Networks (CNNs) yielded 79.8% accuracy, the highest of the algorithms we tested, there's still room for improvement, and we believe that further research in this subject should be done.

Contents

1 Introduction

It is a fact that racism exists in the health industry and affects billions of people around the world. It manifests as a systematic discrimination in countless facets and aspects of the health industry to the point of being supported through institutional policies. This means that it is not practiced through the aberrant behavior of a handful of individuals, although this remains possible, but through unintentional and unconscious bias based on negative stereotypes, and is supported institutional policies. [23]

One noticeable example of this is the significant difference in infant mortality rates and pregnancy related deaths between various races. Between 1999-2013, infant mortality rates were highest among infants born to non-Hispanic black women with 11.11 infant death for ever 1,000 live births. Infants born to non-Hispanic Asian or Pacific Islander mothers experienced the lowest rates of 3.90. In 2014, non-Hispanic black mothers also had the highest percentage of preterm births of the five racial and ethnic groups evaluated in said study (non-Hispanic blacks, non-Hispanic American Indians or Alaska natives, Hispanic or Latina, non-Hispanic whites, and non-Hispanic Asians or pacific islanders) [2]. Black women were also found to be three to four times more likely to die in pregnancy related births when compared to white women. [3]

Another example lies in the distribution of primary care physicians by zip code. Gaskins *et al.* found that the odds of a given zip code having a primary care physician shortage were 67 percent higher for majority African American zip codes but 27 percent lower for majority Hispanic zip codes. This number varied with the amount of segregation in said area: as segregation increased, so did the odds of having a shortage of primary care physicians. The inverse was true for primarily Hispanic and Asian zip codes. [13]

One more example lies in the astonishing rate that minority medical patients are disbelieved or not taken seriously about their health. When polled about perceived pain in white and black patients, a study found that a substantial number of white laypeople and medical students and residents held false beliefs about biological differences between black and white patients. This phenomenon is particularly dangerous because it can cloud a health care provider's judgment and effect their treatment recommendations and diagnosis accuracy. [15] In addition, Asian patients were found to be 1.21 times more likely than white patients to be admitted to the hospital following an ED visit and Black patients were 7% less likely to receive an urgent ESI

score than white patients to immediate or emergent ESI scores, as opposed to semi- or non-urgent scores. Black patients were additionally 10% less likely than white patients to be admitted to a hospital and were 1.26 times more likely than white patients to die in the ED or hospital. [25] While many of these studies were based in the United States, it is evident that this problem is international because Public Health England found that People of Chinese, Indian, Pakistani, Other Asian, Caribbean and Other Black ethnicity had between 10 and 50% higher risk of death from the COVID-19 when compared to White British, which mirrors a similar theme in the United States. [12]

In addition to the statistics and examples listed here, one should wonder what additional ways has racism affected people that have been covered up, or simply not studied yet.

On a separate note, melanoma is a type of skin cancer that presents itself initially as a marking or deformity on the skin. In the year 2020, invasive melanoma is estimated to account for 1% of all skin cancer cases as well as the vast majority of skin cancer deaths. It is also estimated to be one of the most common cancer diagnoses for Americans aged 20-29 years old.[21]

Racism can apply to the already deadly melanoma cancer in a number of ways Despite the occurrence of melanoma in non-Hispanic Black people, survival rates for them lag behind that if non-Hispanic White populations. [10, 11] Skin cancers in people of color often present atypically or in advanced stages when compared to Caucasian patients at the time of diagnosis, with post-metastization diagnoses at rates of 15% of Hispanics, 13% of Asians, and 12% of Black patients, compared with six percent of non-Hispanic Whites for men and seven percent of Hispanics, 21% of Asians, and 19% of Black patients were diagnosed with late-stage melanoma compared with four percent of non-Hispanic Whites. [14, 9] Out of patients who received surgical treatment for melanoma-related complications. The 10-year melanoma-specific survival was 73% lower in black patients than in white patients and other races who had survival rates of 88% and 85% respectively. [8]

To solve the intersection of racism and its effects on melanoma patients, we propose using Artificial Intelligence algorithms (AI) and machine learning as a tool for detecting melanoma in its early stages. They have become increasingly more accessible due to recent advances in access to large data sets, fast computing, and cheap data storage, which has encouraged the development of machine learning algorithms with human-like intelligence in dermatology. [7] However, to prevent hampering this tool's effectiveness and

producing ill effects similar to that stated in previous findings, it is crucial that these algorithms are built to be free from racial bias, especially since melanoma presents on the skin. This can be done on training an AI based on melanoma cases with a variety of skin tones, equalizing these skin tones, and taking into account the relative sparsity of melanoma cases among minority patients.

2 Methodology

Below we outline the various pre-processing and classification algorithms we considered for building a model to detect melanoma in images.

2.1 Pre-Processing Algorithms

2.1.1 CLAHE

Contrast-limited adaptive histogram equalization (CLAHE) is the current state-of-the-art method used for histogram equalization that addresses a lot of the issues other algorithms struggle with [6]. Normal histogram equalization involves distributing the contrast values of the entire image evenly across the entire contrast spectrum. While this works for a uniform condensed distribution of contrast of the image, when there exists varying distributions of contrast across different sections of the image, normal histogram equalization fails to sufficiently equalize all sections of the image. This is where the adaptive part of the CLAHE algorithm is important, where instead of balancing color contrast based on the entire image, CLAHE breaks the image up into tiles that are subsections of the image. These tiles are then used to perform histogram equalizations on subsections of the image. Due to separating the image into smaller regions however, adaptive histogram equalization (AHE) algorithms suffer from a new problem of being more prone to noise in the image [19]. This is where the contrast-limited portion of the CLAHE method is useful. In order to limit the over-amplification of contrast due to noise, CLAHE defines a threshold value for which only the highest parts of the histogram are clipped off and then equalized [1].

The CLAHE algorithm requires 3 input parameters to operate on. The first one being the input image. This is generally a black and white image, however the CLAHE method can also be applied to colored images as well. In the case of colored images, instead of equalizing each of the 3 red, green, and blue spectrums, it is usually recommended to convert the image to the HSV color space and only apply CLAHE to the luminance channel [24]. The other two inputs that were described above are the clip limit value and the grid size for the tiles. The CLAHE algorithm will then run the operations outlined above on the imputed image and output an image with equalized contrast.

2.1.2 Histogram Equalization in MatLAB

There are a variety of methods to accomplish Histogram Equalization in Matlab. The first of which includes the MatLAB Specific command, Histeq. [18] This algorithm is particularly useful because it's easy to use and it works with little to no modifications, allowing the algorithm to be rather short. The command that this algorithm is based on is standard with Matlab's Image Processing Toolbox, and even allows users to input a target histogram should the user want to limit the number of "bins" produced. Despite these benefits, there are a number of caveats to using this algorithm. It can only work with one color channel at a time, so this algorithm would need be run on the red, green, and blue color channels of an image, and then the resulting three histograms could be composited into one equalized image. Unfortunately, this algorithm is difficult to adjust should the user want to alter the method used. This can be solved by building a histogram equalization method from scratch.[4] Dr. Agrawal's algorithm is especially useful because it doesn't use the Matlab specific histogram equalization command, a user can adjust the algorithm as much as is necessary to achieve the desired result. This, of course, will require more expertise in histogram design and equalization methods, but this is worthwhile for the benefits of this algorithm.

2.1.3 LB-CLAHE

LB-CLASH is a machine learning based version of CLASH, it seems to have slightly better results than the standard version of CLASH, but there is currently only one article about the topic (1). Machine learning was used in the artial to control and optimize the hyperparameters for CLASH histogram equalization. The Authors found that LB-CLASH was able to quickly suggest the hyper parameters for a CLASH based histogram equalization that provide better image results. Machine learning does not do the histogram equalization; it just suggests parameters. The article asserts that this method of selecting the hyper parameters is more effective than the other options. They do provide the source code for this project (2), along with the datasets that they used. This method may be usable for our purposes but it will increase the time to compute because before CLASH can be run it must first have its hyper parameters have been determined by LB-CLASH.

2.2 Classification Algorithms

2.2.1 Convolutional Neural Networks (CNNs)

Convolutional Neural Networks (CNNs) are used to classify data based on their internal knowledge base acquired from training on past data. CNNs as a whole is a large topic, so this writeup will mainly focus on the aspects that are relevant to image classification. Functionally, a CNN is a structure of nodes similar to Artificial Neural Networks (ANNs), so it is important to understand how ANNs work before diving into CNNs. ANNs have 3 layers of nodes, the first being the input layer where each predictor for the classification is inputted into the system. ANNs then have an output layer and some number of hidden layers containing perceptrons. Each of these perceptrons have an activation function where, depending on the input values and their corresponding weights, the perceptron outputs a corresponding classification. These perceptrons together as a network output an overall classification for the data from the output layer. The reason why training data is required for neural networks is because initially the weights for how much each predictor or input value should be considered is unknown. So by showing the model a large sample size of pre-classified training data allows it to determine the optimal weights for each predictor or input value in order to achieve the correct classifications.

CNNs depend on this fundamental network of perceptrons, however they take it a step further with convolutional layers. While ANNs can technically be trained to perform image processing and classification problems, it would require an unwieldy amount of weights and layers as every pixel in the image would have to be an input. With convolutional layers however, the amount of required parameters is significantly reduced as the image is split up into tiles. The size of the tile is defined by the model's batch size [16].

In summary, the notable hyperparameters for initializing a CNN model include: number of convolutional layers, initial weights, type of activation function, number of epochs (number of iterations the model trains on the data), and batch size [20]. Once the model has been defined, it is then trained and tested on a pre-classified database until a sufficient success rate is achieved. Finally once the model is trained, the only step required to use it is to input a given image, the CNN model will process the image, and then the model will output a corresponding classification with a given degree of certainty.

2.2.2 Edge Detection

MatLAB offers a selection of edge detection algorithms, which may be useful in detecting the boundaries of an instance of Melanoma.[17] t relies on discontinuities in brightness to detect an "edge". There are several algorithms that can be used to detect these edges. The examples given use both the Sobel and Canny methods This algorithm has demonstrated promising results on very clearly defined melanoma occurrences. Further testing of this algorithm is required for further confirmation on this. However, this algorithm has demonstrated weaknesses when used on images with textured skin. Applying a blurring filter to an image can improve on this weakness, but a better system to remove skin texture is necessary for this algorithm to improve. Further investigation into which edge detection algorithms work best for certain instances of melanoma certainly deserve further investigation.

2.2.3 Haar Cascade

Haar Cascade is a machine learning object detection algorithm conceptualized by Paul Viola and Michael Jones in 2001 The algorithm is used to identify objects, specifically faces and body parts, however it can be implemented to identify many other objects, including melanoma in skin cancer patients. There are four stages to the algorithm: Haar Feature Selection, Creating Integral Images, Adaboost Training, and Cascading Classifiers. Step one in the algorithm is to collect the haar features, which are adjacent rectangular regions at a specific location in a detection window, using Integral Images makes feature collection fast and easy. The next step is to implement the Adaptive Boosting meta algorithm to make sure that the classifiers select the best features available. Finally, in the Cascading Classifiers, each stage consists of decision dumps that decide whether or not there is an object within the window that they're analysing. If there is an object then the detector will mark it as a positive detection, if there is no object then the detector will mark it as negative. Sometimes a detector will incorrectly detect a positive or negative object. In order to combat this multiple stages are implemented to properly filter out false positives and negatives.

To test the algorithm the user must give the program a large quantity of positive and negative objects, which would be images of patients with and without melanoma in order to "train" the algorithm to know what to look for when analyzing images. This algorithm may have a hard time properly

detecting melanoma in skin cancer patients because it can look very different from patient to patient, however, if enough images are given this algorithm could be very promising in its ability to "learn" how to detect a given object. If it is given enough images of positive and negative images, like images of moles and malignant tumors and other growths, this algorithm would have no trouble properly detecting melanoma in skin cancer patients.

2.2.4 Histogram of Oriented Gradients

Histogram of Oriented Gradients(HOG) is an algorithm founded by Robert K McConnell in 1986 but was popularized by French researchers Navneet Dalal and Bill Triggs. The algorithm is designed to take an image, filter it, and categorize it by its distribution of directions of gradients. In order to do this, HOG first reshapes all the images to be the same dimensions and normalizes them to get rid of any possible illumination effects. Next the gradient images are calculated by finding the magnitude and direction of each gradient, using g=gx2+gy2and=arctan(gy gx), respectively. The next step is to calculate the histogram of the gradients using 8x8 cells. Afterwards is the 16 x 16 block normalization which is done by concatenating 4 8x8 cells into a 36 x 1 vector and dividing each element in the vector by its length. The length of a vector is calculated by squaring each element in the vector and adding them together and then finding the square root of that sum. Finally, the HOG feature vector is calculated by concatenating 105 positions of 36x1 vectors together to get a 3780 dimensional vector.

In order to run the algorithm, an image must be given. Based on the images given the code will resize the image to be 32x32 if it isn't already, puts it in greyscale, then a histogram is calculated and a 4x4 matrix is designed for each cell. This algorithm could be useful towards accurately pointing out what patient has melanoma and what patient does not. Since melanoma has distinguished features they could be picked up by the algorithm and accurately categorized in a bin. This algorithm can also be "trained" on finding specific features so if enough images are put into the algorithm they could be accurately categorized based on if they meet the criteria for melanoma.

2.2.5 Support Vector Machine

Support Vector Machine (SVM) is a supervised statistical learning algorithm that is used in classification in machine learning. The main component of

an SVM is the hyperplane, a line that serves as a decision boundary for the n number of features in each data set. The goal of SVMs is to widen the decision boundary such that future datapoints can nearly be categorized in a chosen class. The boundary can be transformed by manipulating the hyperparameters of the algorithm, such as: C (the distance of the decision boundary), gamma (the separateness of the features), the kernel (form of the decision boundary), etc. This is an exploration on using SVMs to predict whether an image has melanoma based on this paper from the Journal of Clinical Oncology. This section will detail the methodology used to implement this algorithm and the results of running the SVM on the data set. The C value had to be fine-tuned and the overall grid had to be optimized to increase the accuracy rate. The gamma and the kernel were left unchanged. The SVM was tested to use three different arrays of C values: 1) 0.1 – 1000, 2) 0.5 – 5000, and 3) 1 – 10000. Once the best SVM C values was found, the code was re-run to get baseline accuracy scores. Afterwards, the grid-search was optimized specifically to improve accuracy scores.

3 Dataset

We had an initial melanoma dataset provided to us for training and testing our models, however we eventually realized we needed a much larger dataset to sufficiently train our classification algorithms. Ideally, we wanted a dataset that both contained an even distribution of labeled melanoma and non-melanoma images as well as an even distribution of labeled varying skin tones in the images. This would allow us to not only determine our model's overall accuracy at detecting if melanoma was present in an image, it would also allow us to determine how fair our model predicted across various skin tones. Unfortunately, we were unable to find a dataset with the latter labels of skin tones. We did, however, found a relatively large dataset, "Human Against Machine with 10000 training images" (HAM10000) [22] containing labeled data of a variety of different pigmented lesions. From this dataset, we extracted all of the melanoma labeled images as well as an equal number of randomly selected non-melanoma images. This gave us a total of 1192 images, which we then shuffled and randomly split into roughly 80% training and 20% testing data.

4 Implementation and Results

For testing all of our preprocessing and classification algorithms, we designed an all-in-one framework that allowed us to dynamically run any classification algorithm with a given set up hyper-parameters to optimize over, and a given set of pre-processing algorithms to run on the input data beforehand. As a result, we were able to both easily run or swap out different classification or pre-processing algorithms to compare different combinations. We also built this framework in Google Collab, allowing our entire team to easily share, edit, and run each other's code on Google Collab's fast GPU. Note: Not all algorithms from the Methodologies were implemented due to knowledge and time constraints, however in future works these could be areas to further explore.

4.1 CNN Classification with Global HE and CLAHE

When designing a Convolutional Neural Network (CNN) for classifying our melanoma dataset, we decided to replicate the general structure of other popular and proven successful CNN architectures. Specifically, the model consisted of 5 convolutional layers which is the standard for extracting key spatial features in the image, as well as 2 fully connected layers used for determining relationships between the found features and the target classifications. This setup is similar to that of CNN models VGG16 and AlexNet, which are ILSVRC challenge winners in 2014 and 2015 respectively. [5]

Getting more in depth with the individual layers, the input shape of the first convolution layer was set as a standard 224x224x3 colored image. As a result, we had to reshape all the data to this size by bilinearly resizing the images in order to best preserve the same image structure. For each convolutional layer we used 64 filters of 3x3 kernels and a stride of 1. After each convolutional layer we added 2x2 max pooling layers which helps the model generalize the features and as a result decreases overfitting on the training data. For the fully connected layers, the first layer had 4096 nodes and the second layer had 2048 nodes, with both using relu activation functions. After each dense layer we added dropout layers which randomly drop nodes during each batch in order to ensure all nodes are being trained, giving the model more expressive capabilities with more influential nodes. Finally, the output layer consists of 2 nodes, corresponding to the two classifications of melanoma or not melanoma, as well as a softmax activation to select the higher of the

two classifications. For the model's hyperparameters, after optimizing the model over a gridsearch of parameters, we obtained our best model with a batch size of 16, epochs of 50, loss function of categorical cross entropy, and an Adam optimizer.

Another important aspect of our initial research was trying to find preprocessing algorithms that would improve our classification algorithm's results. We decided that the global histogram equalization as well as the local histogram equalization, specifically CLAHE, methods would be applicable to input data of our CNN. Predictably, applying CLAHE to the input data achieved better results than applying global histogram equalization, with CLAHE yielding a 72.7% accuracy and global HE yielding only a 62.6% accuracy. Surprisingly however, without any pre-processing our CNN achieved the highest total test accuracy of 79.8%. We had initially theorized that applying a histogram equalization method would help decrease the differences in skin tones in order to increase the fairness of the model's accuracy between races, which in turn would increase the model's overall accuracy. While we didn't have access to labeled skin tones for our data to determine whether this preprocessing evened the model's accuracy across races, we can determine that overall, this preprocessing had additional side effects that lowered the overall accuracy of the model.

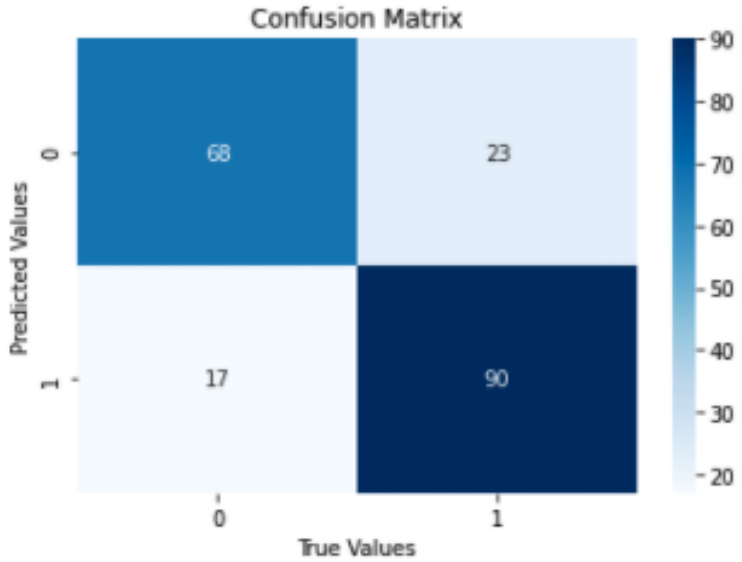

Figure 1: CNN Confusion Matrix with No Preprocessing

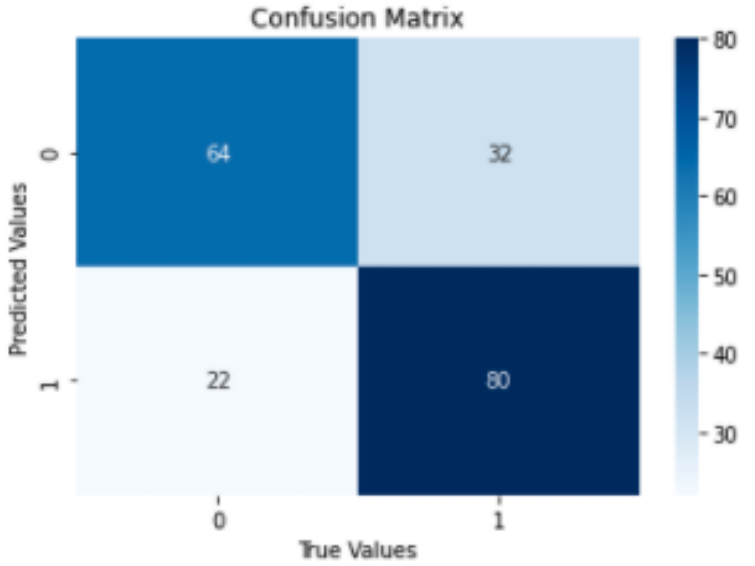

Figure 2: CNN Confusion Matrix with CLAHE Preprocessing

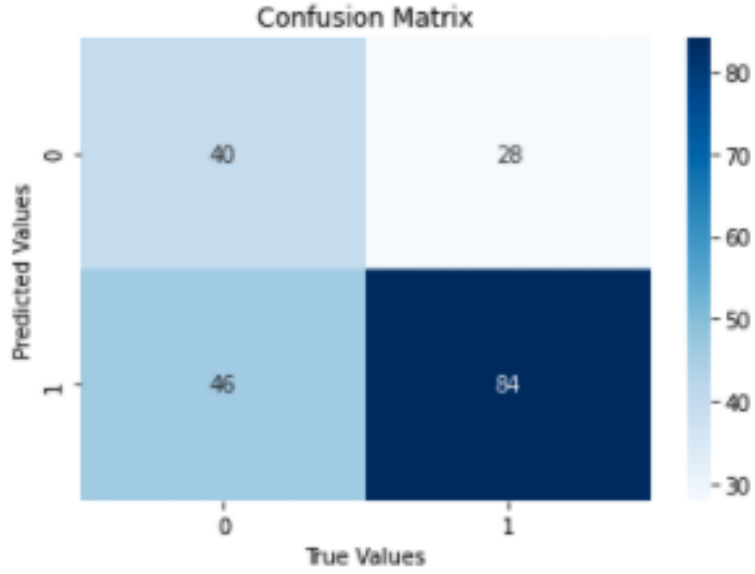

Figure 3: CNN Confusion Matrix with Global Histogram Equalization Preprocessing

4.2 Histogram of Oriented Gradients (HoG)

Histogram of Oriented Gradients (HOG), focuses on the structure of objects. It extracts information about edges including their magnitude and orientation. This alone is not an image classifier, it is more of a pre-pressing. By hooking the features or gradients extracted by HOG into a machine learning framework we can then use those features to attempt image classification.

To extract the HOG features from each image the python library skimage was used. Each image was processed with 16x16 Pixel cells and 1 cell per block. Those features were then passed into a fully connected neural network created using the tensorflow python library. The was created with two relu layers and one softmax regression layer. The machine was then trained and the results are shown below.

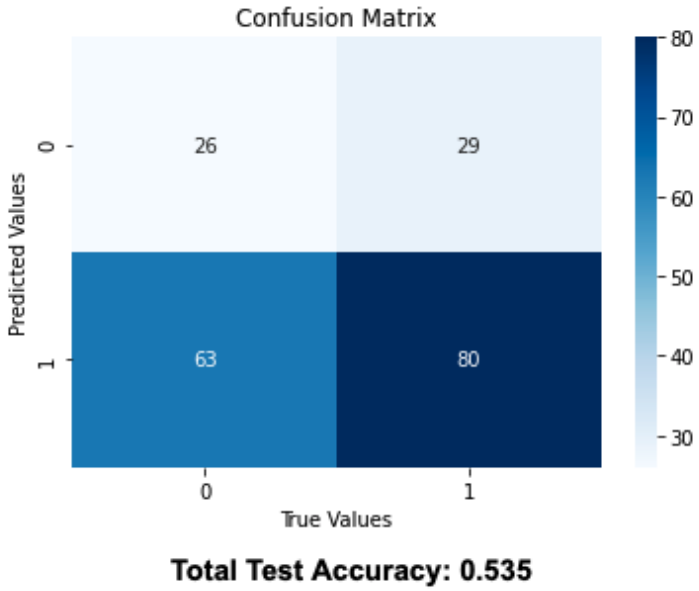

Figure 4: HoG Confusion Matrix

With the best test accuracy of 53% being very little better than just guessing, it can be inferred that using HOG as an approach to classify melanoma is not possible. This may be due to the method's shortcomings in identifying fine details. HOG is a method that Eccles with edge detection tasks and general shapes, so something, as finessed as melanoma, is difficult for this type of classification.

There is other possibilities, the training and testing data sets were not great and could have lead to the failure of the algorithm. Although if a better dataset was used the results would likely not change much, due to the fact that HOG is not good at classifying things other than images.

4.3 Haar Cascade

To create a Haar Cascade classifier for identifying cancerous melanoma we had to prepare a number of samples, which would then be used to train the algorithm provided by the OpenCV python library. First a batch of negative samples was collected from our dataset, which would be used to identify background noise we were not interested in. Then a batch of positive samples containing cancerous melanoma was selected. However, the setup for the algorithm also required annotations for the positive samples, which we created using a tool provided by OpenCV. These annotations must indicate how many instances of the object we're interested in appear in the image, where they're located and their size.

We were able to prepare 100 negative samples and 60 positive samples. The model was also set to have 15 stages, each one being more precise but also more time-consuming, causing only promising candidates to be closely inspected and obvious false cases to be quickly discarded. Once the model was trained, we moved on to test the accuracy of its predictions.

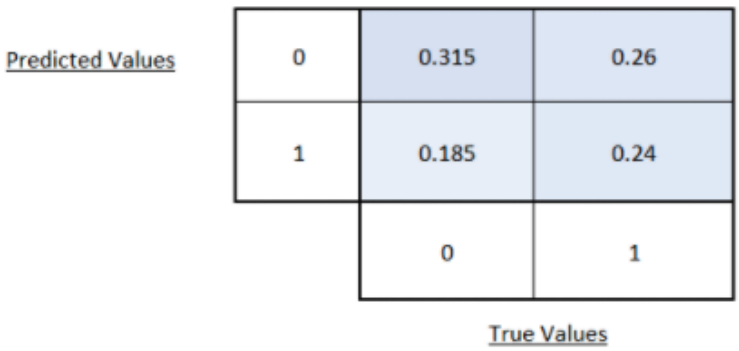

Figure 5: Haar Cascade Confusion Matrix

The finished model had an accuracy rate of 55.5%. These results are not very promising, with the model being only slightly better than random

guesses. One reason why the model may have underperformed is the lack of samples. Since each positive sample must be annotated by hand, this presents a major limitation on how many images can be used for training and testing. Time constraints and lack of manpower meant that, even though we had collected many samples, we were only able to use a small subset of them to train our haar cascade classifier.

However, we believe this algorithm is not the right tool for this particular kind of job. Haar cascades algorithms are generally used to quickly detect whether an object of interest is present in an image, but not to detect subtle differences between two similar kinds of object. Time is not a critical factor, so the ability of the haar cascade classifier to quickly work through images is wasted. Furthermore, the annotation tool allows for multiple instances of an object to be detected by the classifier, but this is never taken advantage of since each of our samples contains one skin pattern each. Because of this, we don't believe the haar cascade classifier will provide reliable results, even if more samples were collected.

4.4 Support Vector Machine

The results of the SVM is that it achieved 71% accuracy in differentiating between melanoma and otherwise. In a testing sample of 7, it was able to properly predict 5 of 7 images correctly.

	precision	recall	f1-score	support
0	0.50	1.00	0.67	2
1	1.00	0.60	0.75	5
accuracy			0.71	7
macro avg	0.75	0.80	0.71	7
weighted avg	0.86	0.71	0.73	7

```
[[2 0]
 [2 3]]
```

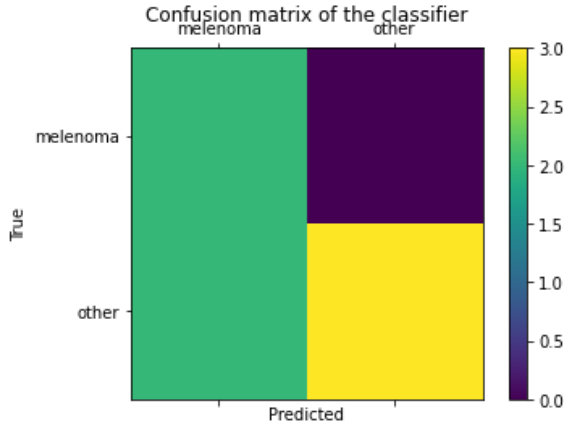

Figure 6: Support Vector Machine Matrix

5 Conclusion

Surprisingly, our CNN algorithm on its own performed best, even when unassisted by CLAHE and HE. While, a 79.8% accuracy rate is a promising start, there is room for improvement. Due to the limited number of algorithms that we were able to research and test, it is very likely that more extensive research will yield more convincing results and find tools better suited to the finding of cancerous melanoma. It's possible that some of the algorithms employed were not the best tool for detecting melanoma, or that their training process can be refined to yield better results, even with no other significant changes. Regardless, we believe this venture is worth investigating further.

Table 1: Table showing accuracy of each algorithm.

Algorithm	**Accuracy**
CNN	**79.8%**
CNN & CLAHE	72.7%
CNN & HE	62.6%
HoG	53.5%
Haar	55.5%

References

[1] Histograms - 2: Histogram equalization. *OpenCV*.

[2] *Health, U.S., 2015: With Special Feature on Racial and Ethnic Health Disparities;2016 ASI 4144-11;DHHS Pub. No. 2016-1232*. 2016.

[3] Investigators at women's health research institute report findings in obstetrics and gynecology (reducing disparities in severe maternal morbidity and mortality). *Women's health weekly*, pages 193–, 2018.

[4] Rashi Agrawal. Implement histogram equalisation without histeq( ), Aug 2014.

[5] Emna Amor. 4 cnn networks every machine learning engineer should know. *TOPBOTS*, 2020.

[6] M. Bhat and M. S. Tarun Patil. Adaptive clip limit for contrast limited adaptive histogram equalization (clahe) of medical images using least mean square algorithm. *2014 IEEE International Conference on Advanced Communications, Control and Computing Technologies*, pages 1259–1263, 2014.

[7] Stephanie Chan, Vidhatha Reddy, Bridget Myers, Quinn Thibodeaux, Nicholas Brownstone, and Wilson Liao. Machine learning in dermatology: Current applications, opportunities, and limitations. *Dermatology and therapy*, 10(3):365–386, 2020.

[8] Karen Kadela Collins, Ryan C Fields, Dadrie Baptiste, Ying Liu, Jeffrey Moley, and Donna B Jeffe. Racial differences in survival after surgical treatment for melanoma. *Annals of Surgical Oncology*, 18(10):2925–2936, 2011.

[9] Rosemary D. Cress and Elizabeth A. Holly. Incidence of cutaneous melanoma among non-hispanic whites, hispanics, asians, and blacks: An analysis of california cancer registry data, 1988-93. *Cancer causes & control*, 8(2):246–252, 1997.

[10] MaryBeth B Culp and Natasha Buchanan Lunsford. Melanoma among non-hispanic black americans. *Preventing chronic disease*, 16:E79–, 2019.

[11] Sean M. Dawes, Sheena Tsai, Haley Gittleman, Jill S. Barnholtz-Sloan, and Jeremy S. Bordeaux. Racial disparities in melanoma survival. *Journal of the American Academy of Dermatology*, 75(5):983 – 991, 2016.

[12] Public Health England. Disparities in the risk and outcomes of covid-19. 2020.

[13] Darrell J Gaskin, Gniesha Y Dinwiddie, Kitty S Chan, and Rachael R McCleary. Residential segregation and the availability of primary care physicians. *Health services research*, 47(6):2353–2376, 2012.

[14] Hugh M Gloster and Kenneth Neal. Skin cancer in skin of color. *Journal of the American Academy of Dermatology*, 55(5):741–760, 2006.

[15] Kelly M. Hoffman, Sophie Trawalter, Jordan R. Axt, and M. Norman Oliver. Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites. *Proceedings of the National Academy of Sciences*, 113(16):4296–4301, 2016.

[16] Abhig Master. Difference between ann, cnn and rnn. *GeeksforGeeks*, Jul 2020.

[17] MathWorks. edge, 2006.

[18] MathWorks. Ihisteq, 2006.

[19] P. Musa, Farid Al Rafi, and Missa Lamsani. A review: Contrast-limited adaptive histogram equalization (clahe) methods to help the application of face recognition. *2018 Third International Conference on Informatics and Computing (ICIC)*, pages 1–6, 2018.

[20] Pranoy Radhakrishnan. What are hyperparameters ? and how to tune the hyperparameters in a deep neural network? *Medium*, Oct 2017.

[21] Rebecca L Siegel, Kimberly D Miller, and Ahmedin Jemal. Cancer statistics, 2020. *CA: a cancer journal for clinicians*, 70(1):7–30, 2020.

[22] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5:180161, 2018.

[23] D.R Williams and T.D Rucker. Understanding and addressing racial disparities in health care. *Health care financing review*, 21(4):75–90, 2000.

[24] Wan Yussof. Performing contrast limited adaptive histogram equalization technique on combined color models for underwater image enhancement. 01 2013.

[25] Xingyu Zhang, Maria Carabello, Tyler Hill, Sue Anne Bell, Rob Stephenson, and Prashant Mahajan. Trends of racial/ethnic differences in emergency department care outcomes among adults in the united states from 2005 to 2016. *Frontiers in medicine*, 7, 2020.