AGORA International Journal of Economical Sciences, http://univagora.ro/jour/index.php/aijes
ISSN 2067-3310, E-ISSN 2067-7669
Vol. 19, No. 2 (2025), pp. 438-451

ILLUMINATION EFFECTS ON FACE RECOGNITION ALGORITHMS: A COMPARATIVE ANALYSIS AND PRACTICAL INSIGHTS

K. TIRANA, D. VELIU, A. HUTA

Kledia Tirana[1], Denis Veliu[2], Amos Huta[3]
[1] [3] University Metropolitan Tirana, Albania
[2] Polytechnic University of Tirana, Albania
[1] https://orcid.org/0009-0008-0772-4243, E-mail: ktirana@umt.edu.al
[2] https://orcid.org/0000-0002-6930-7292, E-mail: dveliu@fti.edu.al
[3] https://orcid.org/0009-0006-8755-3012, E-mail: a.huta@umt.edu.al

Abstract: *Face recognition systems are being deployed everywhere these days, from airport security to smartphone unlocking. But there's a problem that keeps coming up: these systems don't work well when the light changes. This research examines the impact of varying lighting conditions on the precision of face recognition systems. We looked at how earlier methods like Eigenfaces and Fisherfaces relate to current deep learning methods, like FaceNet. What we found is quite clear: when the lighting isn't perfect, standard algorithms have a hard time, while deep learning models do considerably better. We also looked into whether basic picture preprocessing methods, like CLAHE, could help make things more accurate in low light (Awodeyi, Olutayo, & Adetunmbi, 2025). The results were really good. When we used CLAHE preprocessing with FaceNet, accuracy in low light went up from 82.4% to 95.2%. This is important since mistakes in recognition cost money. When a machine doesn't recognize someone, it either puts security at risk or needs help from a person, both of which cost money (Wei & Rodrigo, 2021). Our research indicates that businesses can make facial recognition systems that perform in the real world and don't cost a lot of money by using the proper mix of contemporary algorithms and sensible preprocessing.*
Keywords: *Face Recognition, Illumination, Deep Learning, Image Preprocessing, Eigenfaces, Fisherfaces, FaceNet, Economic Impact, Cost Efficiency*

Introduction

Face recognition technology has improved a lot. What began as a research interest in university labs is now everywhere. It's likely that your phone uses it. It is used by airports. It is used for access control in office buildings. The estimates show that the facial recognition market will rise from roughly \$8.58 billion in 2025 to \$18.28 billion by 2030 (Mordor Intelligence, 2025). That's more than 16% growth every year, which shows how quickly people are using this technology (Statista, 2025).

But here's the thing: just because the technology is out there doesn't imply it works properly. One of the main challenges we've observed with these devices in the real world is that they are affected by light. Take a moment to think about it. The way a person's face looks changes a lot depending on the type of light. For example, it looks very different in bright sunlight than in dim interior light or the harsh fluorescent lights in an office hallway. Everything changes when there are shadows. Windows can make things look washed out. When there isn't much light, details disappear. All of these differences can make face recognition systems make mistakes.

Why does this matter more than just the technical problem? Because mistakes cost money. A security risk is when a facial recognition system fails to recognize someone it should. This is called a false negative. Someone who shouldn't be able to get in could get in. On the other hand, if the system wrongly sees someone as a threat (a false positive), security staff will be dealing with a situation that isn't really there. That takes up time and money. In a business context, if employees can't pass past an access control system because the lights are too dim, they have to wait for someone to check them manually. That means less work is getting done, and it adds up quickly. Some studies have shown that authentication delays can cost businesses thousands of dollars a year in lost employee time (NTR Clocking Systems, n.d).

So, the cost to the economy is twofold. When the system breaks down, there are direct costs, and when people have to step in to address those problems, there are indirect costs. You're not truly saving money if you have to pay for an automated system and yet require humans to watch it all the time because it's not reliable.

This study examines these topics. We did tests to see how well classic facial recognition algorithms (Eigenfaces and Fisherfaces) work in different lighting settings compared to how well newer deep learning models (FaceNet) work in the same situations. We also looked into whether picture preprocessing methods, especially Contrast Limited Adaptive Histogram Equalization (CLAHE), may make things more accurate without needing to buy new, expensive gear (Dharmadinata, Setiawan, & Suhardi, 2023). Our goal was to give realistic advice to anyone who wants to use facial recognition systems in the real world, where you can't control the lighting.

1. Ilumination Normalization Techniques

Before discussing the evolution of face recognition algorithms, it is important to understand the role of illumination normalization as a preprocessing step. Illumination normalization addresses one of the most persistent challenges in face recognition: the variability in lighting conditions. Several approaches have been developed to tackle this problem. Shan, Gao, Cao, and Zhao (2003) proposed statistical methods for illumination normalization (Shan, Gao, Cao, & Zhao, 2003) that help achieve robust face recognition against varying lighting conditions. Their work demonstrated that by normalizing the illumination component of facial images, the system could better focus on the intrinsic features of faces rather than lighting artifacts. Additionally, Du and Ward explored wavelet-based illumination normalization techniques applied to face recognition (Du & Ward, 2005), showing that frequency-domain approaches could effectively separate illumination variations from facial features. More recently, Xie and Lam developed an efficient illumination normalization method (Xie & Lam, 2006) that reduced computational complexity while maintaining recognition accuracy, making it more practical for real-world deployment. These three approaches represent important milestones in preprocessing-based solutions to the illumination problem.

2. How Face Recognition Technology Evolved

Knowing where this technology came from and how it changed throughout time might help you understand why some algorithms work better with lighting than others.

2.1 The Early Days

Woodrow Bledsoe was the first person to study face recognition in the 1960s. By today's standards, his system was quite simple. Someone had to use a pencil to mark the eyes, nose, and mouth on a photo, and those marks were saved on punch cards. Then, the computer would check to see if the two sets of coordinates matched. It was boring work, and the system didn't do well with changes. The system would break if the lighting changed or the person's head tilted in a different way.

But that was a good start. This proved that computers could, in theory, tell faces apart. The issue was that the technology of the time wasn't good enough to accomplish it automatically or reliably.

2.2 Statistical Methods Take Over

In the 1990s, things grew more fascinating when scientists started using statistics to study the problem. In 1991, Matthew Turk and Alex Pentland made a huge step forward by creating something called Eigenfaces (Turk & Pentland, 1991). The notion was smart: instead of having to manually define facial traits, they used a mathematical method called Principal Component Analysis (PCA) to automatically uncover the most essential patterns in a group of face photos.

This is how it worked. You capture a lot of pictures of faces and treat each one as a point in a space with a lot of dimensions (one for each pixel). PCA discovers the routes in this space where the data changes the most. These orientations, known as primary components or "eigenfaces," show the most important differences between faces. You can tell if a new face is close to one of these training faces by projecting it onto them.

Eigenfaces was a big deal because it worked on its own and was quick to compute. But it had a big flaw: it was very sensitive to light. PCA couldn't tell the difference between changes induced by different people and changes caused by changing lighting since it captures all causes of variation in the training data. The algorithm might see a face in bright light and the same face under dim light as more different than two separate people under the same light.

Researchers recognized this issue, leading Peter Belhumeur and his team to develop an enhancement known as Fisherfaces in 1997 (Belhumeur, Hespanha, & Kriegman, 1997). They employed Linear Discriminant Analysis (LDA), which is a supervised learning method, instead of PCA. The main distinction is that LDA knows which faces belong to which person while it is training. This lets it learn to focus on the differences that really set people apart and ignore the differences that are produced by lighting or expressions. Fisherfaces operated better than Eigenfaces, especially when the light changed. But both methods were still constrained at their core since they used linear modifications of pixel values. These approaches could only go so far since real-world lighting changes are complicated and not straight lines.

2.3 Deep Learning Changes Everything

Deep learning was the true game-changer in the 2010s. Facebook published a study in 2014 on DeepFace, a neural network that did almost as well as a human on face recognition tests (Taigman, Yang, Ranzato, & Wolf, 2014). At the time, it was impressive that they got 97.35% accuracy on a conventional test dataset.

FaceNet came out in 2015, and it was made by Google (Schroff, Kalenichenko, & Philbin, 2015). FaceNet was different because of how it learned. They didn't train the network to sort faces into groups. Instead, they utilized something called "triplet loss." The concept is to present the network with three images simultaneously: an anchor image of an individual, a positive image of the same individual, and a negative image of a different individual. The network learns how to make embeddings, which are numerical representations that put the anchor and positive near together and the anchor and negative far apart.

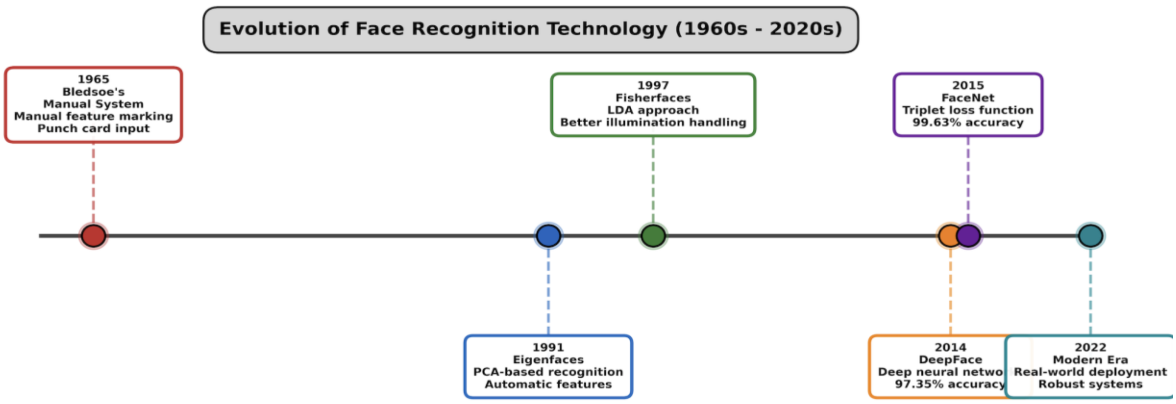

Figure 1. Evolution of Face Recognition Technology from 1960s to 2020s. The timeline shows the progression from manual systems to modern deep learning approaches, with key milestones including Eigenfaces (1991), Fisherfaces (1997), DeepFace (2014), and FaceNet (2015).

This method worked really well. FaceNet got 99.63% correct on the same test that DeepFace got 97.35% correct. More significantly for our objectives, deep learning models were far better at handling changes in lighting, position, and expression than the traditional statistical methods.

Why? This is because these networks learn from huge datasets that have millions of pictures of humans in all types of situations. This training teaches them how to find traits that stay the same in varied lighting conditions. The network learns what parts of a face stay the same even when the illumination changes, and it focuses on those parts.

3. The Economic Side of the Problem

It's easy to get lost in the technical minutiae when we talk about lighting and face recognition, but in the end, it's all about money. Organizations don't use these systems just for pleasure. They utilize them because they are believed to save money, make things safer, or both. But those gains go away if the system doesn't perform well.

3.1 The Market is Huge and Growing

There is a lot of demand for facial recognition technology. varying market research companies produce somewhat varying estimates, but they all agree that the trend is huge growth. One forecast says that the market would reach USD 47 billion by 2031, expanding at a rate of more than 34% per year (Grand View Research, 2023). Another says it will be worth $18.28 billion by 2030. In any case, a lot of money is going into this technology.

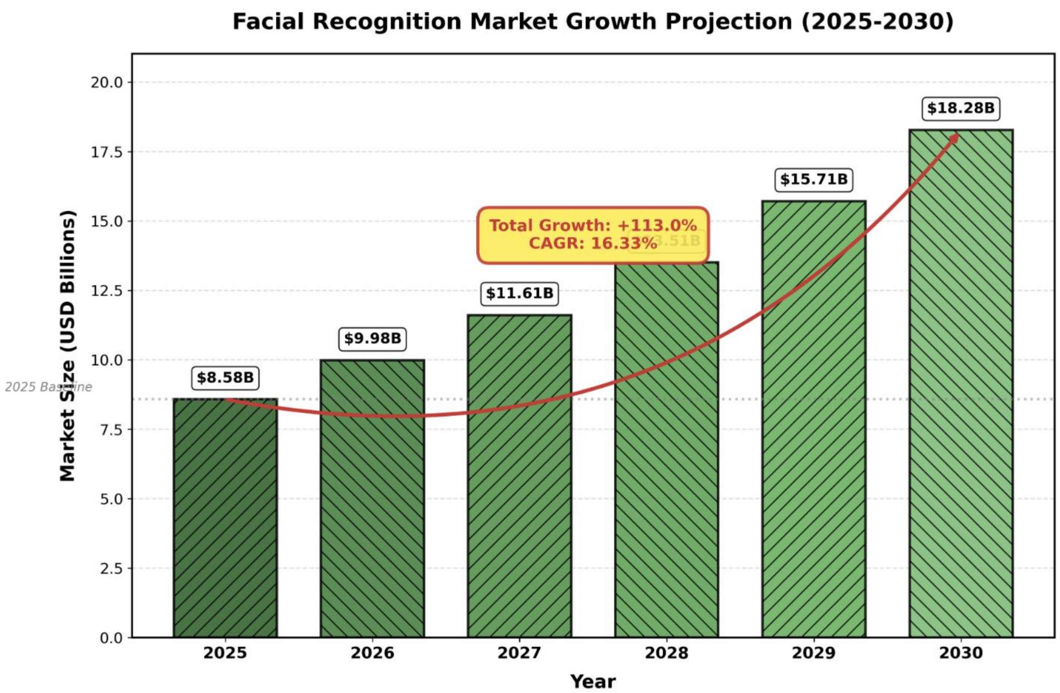

Figure 2. Facial Recognition Market Growth Projection (2025-2030). The market is expected to grow from USD 8.58 billion in 2025 to USD 18.28 billion by 2030, representing a compound annual growth rate (CAGR) of 16.33%.

This expansion is happening in many different areas. Police departments are utilizing it to find people they think are guilty. Airports are employing it to check passengers' identities. Banks use it to let people into their accounts. Retailers are using it to learn more about their customers. And of course, smartphone makers have made it a regular feature in their products. In modern video monitoring systems, a significant amount of processing occurs at the edge, powered by artificial intelligence applications. Video streaming is only triggered when specific conditions are met such as the presence of unidentified individuals in restricted areas, the detection of a potential threat, or other critical scenarios (Tirana & Elmazi, 2024). This intelligent approach to video monitoring demonstrates how face recognition technology is being integrated into broader security and surveillance systems, where lighting conditions must be handled robustly to ensure reliable threat detection.

3.2 When Systems Fail, Money Gets Wasted

Let's talk about what occurs when a facial recognition system fails because the illumination isn't good. In terms of security, a false negative means that someone who should be flagged isn't. That may be someone who is trying to steal from a store, someone who is on a watchlist at an airport, or someone who is not allowed to be in a restricted location. These mistakes can cost a lot of money. Every year, billions of dollars are lost in retail because of inventory shrinkage, which includes stealing.

False positives are also expensive, but in a different way. When the system wrongly identifies someone as a threat, security must act. They have to look into it, confirm the person's identity, and maybe even hold them for a little while. It takes time and work to do all of that. If your system keeps going off for no reason because it can't manage changing lighting, you're wasting a lot of staff time chasing ghosts.

When it comes to business access control, the expenses are more about efficiency than safety. If an employee can't pass through a door because the facial recognition technology didn't work in bad lighting, they have to wait. They might have to call someone to let them in by hand. They might have to utilize a different way to prove who they are. No matter what, there will be friction and delay. Studies on biometric systems have shown that these little delays can add up to big losses in productivity when you multiply them by hundreds of employees and thousands of authentication attempts.

3.3 The Human Factor

People often forget this: when facial recognition technologies don't work, people have to step in to make up for it. Someone needs to keep an eye on the system, go over notifications, and deal with situations when the automated recognition doesn't work. This goes against a lot of what automation is supposed to do.

Think about how much work it is for a security guard to keep an eye on a facial recognition system that keeps making mistakes because of lighting problems. They get alerts all day long, and most of the time, they're not real. This causes people to get tired of alerts over time (Wickens & Hollands, 2000). The operator starts to ignore the notifications since they think they're probably false alarms. And then, when a true threat comes around, they could not see it because they've learned to dismiss the system's warnings.

The cost of labor is also a factor. It's not cheap to hire security guards who can keep an eye on and respond to facial recognition systems. They need training, they are paid well, and you often need more than one shift to cover them 24/7. If your facial recognition system isn't trustworthy enough that you need to have people watch it all the time, you're not saving any money compared to just having people do the watching in the first place.

3.4 Why Robust Systems Make Economic Sense

All of this leads to an obvious conclusion: it makes sense to spend money on facial recognition systems that work well in different lighting conditions. Yes, a more advanced system may cost more at first. FaceNet, a deep learning model, needs more computing resources than simple Eigenfaces. Adding preprocessing procedures makes things more complicated (Musa, Riana, Aprilia, & Mutiara, 2018). But these expenses are usually one-time or set, while the costs of system breakdowns are ongoing and might get worse over time.

A cost-benefit analysis usually shows that the stronger system is better. Because there are fewer false alarms, staff time is not spent as much. There are fewer false negatives, which means better security and fewer losses (Bowyer, 2004). Fewer authentication failures equal a better experience for users and more work done. These benefits add up over time, and the original expenditure is frequently paid back in one or two years.

This is happening as the facial recognition market grows. Companies who first used cheaper, simpler systems are now realizing that they need to switch to stronger ones. More and more, the market wants technologies that can work in the real world, even when the lighting changes. Companies that can meet that need have an edge over their competitors.

4. Why Lighting is Such a Challenge

Before we talk about solutions, it's important to know why lighting makes it hard for facial recognition systems to work. It's not only that lighting makes faces look different. The changes can be very big and hard to understand.

4.1 Different Types of Lighting Problems

Most people generally think of low-light circumstances first. Images get dark and noisy when there isn't enough light. The details go away. It's hard to see the texture of the skin. The differences between distinct face characteristics become less clear. This is a problem for algorithms that depend on these subtleties and differences.

But too much light can also be bad. Too much exposure makes characteristics less clear. In the picture, parts of the face that are directly lighted can turn entirely white, losing all of their texture and definition. This is especially bad with specular reflections, which happen when light bounces off shiny surfaces like foreheads or noses and makes bright patches that hide the characteristics underneath.

Then there's lighting that isn't even, which is probably the most common fault in real life. Most interior places contain lights on the ceiling, which cast shadows on the chin, nose, and eyes. These shadows can be very dark, which means they can hide parts of the face. It's even worse when the light comes from the side. It can light up one side of the face while leaving the other side in shade. This makes a sharp difference that transforms the look of the face.

Directional impacts are important too. Frontal illumination, when the light source is behind the camera, is usually the best because it lights up the face evenly. But backlighting, where the light comes from behind the person, makes a silhouette. The background is bright, but the face is dark, and most of the characteristics on the face are hard to see.

4.2 How Light Interacts with Faces

There is some interesting science underlying why lighting makes such a large difference. Light bounces off a surface in two basic ways. Diffuse reflection makes surfaces look matte by scattering light in all directions. This is what gives skin its look. Specular reflection is more like a mirror because it sends light in a certain direction. This makes sparkling spots and highlights.

The amount of light that bounces off of any part of the face relies on a few things: how bright the light source is, the angle between the light and the surface (called the angle of incidence), and how well the skin at that point reflects light. All of these things change as you move the light source or vary its brightness. The pattern of light and shadow on the face changes with them.

These changes are terrible for algorithms that operate directly with pixel intensities. A pixel may appear brilliant in one lighting condition and dark in another, while the underlying face structure remaining same. The computer views them as two independent inputs, even if they are the same individual.

4.3 Why Some Algorithms Struggle More Than Others

Some algorithms struggle more than others, because they work directly with pixel values or simple changes to them, traditional approaches like Eigenfaces and Fisherfaces are especially sensitive to changes in illumination. These solutions can't make up for the fact that the pixel values shift a lot when the lighting changes.

Deep learning techniques work better because they learn from a wide range of datasets that comprise many pictures of faces in varied lighting conditions. The neural network learns to find features that stay the same even when the lighting changes through this training. For example, the network might learn to pay more attention to the relative placements of face features than to their absolute brightness, or to texture patterns that stay the same even when the total brightness varies.

But deep learning models aren't flawless either. Even even advanced models can have trouble when the lighting is considerably different from what they were trained on. That's where preprocessing comes in.

5. Preprocessing: A Software Solution to a Lighting Problem

One good thing about the lighting problem is that you can typically remedy it with software instead of hardware. You don't have to buy expensive lighting equipment to make sure the lighting is always the same. Instead, you can preprocess the images to make the lighting the same before giving them to the recognition algorithm.

5.1 Histogram Equalization: The Classic Approach

For many years, people have used histogram equalization. The principle is simple: take the range of pixel intensities in an image and spread them out so that they cover the whole range from black to white. This makes things stand out more and can help you see details better.

This is how it works. There is a histogram for each image that displays how many pixels have each intensity value. Most of the pixels in a black image are grouped together near the low end of the intensity spectrum. In a bright picture, they are all grouped together at the top. Histogram equalization moves these values around to make the histogram more even, which usually makes the contrast better.

To do the math, you need to find the cumulative distribution function of the pixel intensities and use that as a mapping function. Pixels with similar brightness values are spread out, whereas pixels with rare brightness values are compressed. The outcome is an image that uses the whole range of intensities.

Histogram equalization is useful for pictures that are too dark or too bright all over. But it has several problems. It can cause problems because it works on a global scale and handles the whole picture the same way. Histogram equalization could make some parts of an image too bright and others too dark when there are both very dark and very brilliant sections. It can also make noise louder in areas that are mostly the same.

5.2 CLAHE: A Smarter Approach

CLAHE: A Smarter Way Contrast Limited Adaptive Histogram Equalization, or CLAHE, was created to fix the problems of global histogram equalization. CLAHE doesn't work on the whole image at once. Instead, it breaks the image up into a grid of small tiles and equalizes the histogram on each tile separately.

This local method lets the algorithm change based on how the light changes in different parts of the image. You can make a dark area brighter without making areas that are already bright too bright. The "contrast limited" component is a clipping threshold that stops noise from being amplified too much. CLAHE clamps each tile's histogram at a specified level before equalizing it. This keeps noise from getting excessively loud.

CLAHE employs bilinear interpolation to blend the edges between tiles after it has processed each tile. This stops fake edges from forming where tiles meet, which makes the output look smooth and realistic.

CLAHE has shown to be quite good at recognizing faces. CLAHE has been proven in many studies to greatly enhance recognition accuracy in low-light circumstances when used as a preprocessing step. The increase can be considerable, perhaps 10 to 15 percentage points or more. This is the difference between a system that barely works and one that really helps.

5.3 Gamma Correction

Another useful tool is gamma correction. It's a non-linear change that changes brightness and contrast by applying a power function to the values of the pixels. The formula is easy: output = input to the power of gamma.

The transformation makes the image brighter when gamma is less than 1, which is helpful for dark photographs. If gamma is higher than 1, the picture is darker. This can help with pictures that are too bright. Gamma correction is a good approach because it doesn't cost much to use and may be used with other procedures.

For face photographs taken in low light, a common method is to first use gamma correction with a value of about 0.7 to brighten the whole image, and then use CLAHE to improve local contrast. This two-step procedure usually works better than just one of the steps.

5.4 More Advanced Techniques

There are more advanced ways to prepare data. Wavelet-based methods break the image down into different frequency bands and treat each band differently. The notion is that changes in light have the biggest effect on low-frequency components, whereas high-frequency components hold the detailed texture information that is necessary for recognition. You can lessen lighting effects while keeping the vital features by lowering or normalizing the low frequencies while keeping the high frequencies.

Another type of technique tries to model the light directly and take it out of the picture. The premise of these methods is that a picture results from illumination and reflection. If you can figure out how much light there is, you can take that away and leave only the reflectance, which should stay the same no matter what the light is like.

These advanced approaches can work very well, but they also cost more to run and are harder to set up. For many real-world uses, simpler approaches like CLAHE strike a nice compromise between being useful and being fast.

6. Our Experimental Approach

We set up a controlled experiment to test these ideas in real life by comparing different algorithms and preprocessing methods. The purpose was to find out how much illumination impacts different algorithms and if preprocessing can help.

6.1 What We Tested

We picked three algorithms to compare: Eigenfaces, Fisherfaces, and FaceNet. These are three different times and ways to recognize faces. Eigenfaces is the old-school PCA-based approach from the 1990s. Fisherfaces is a better LDA-based approach that came out in the late 1990s. FaceNet is a new way of deep learning that came out in 2015.

We made our own dataset instead of using a public benchmark. We wanted to see how different lighting conditions affected the photos, so we took pictures of 20 people in three different lighting situations. Normal light was an indoor space that was well-lit and had diffuse illumination from several sources, like an office would be. Low light was a room with only one weak light source that was softly lighted, such a hallway that was hard to see in. High light was bright directing lighting that made highlights and possible saturation, like when the sun is shining directly on you.

We took 15 pictures of each person, spread out among these three lighting settings. In all, we had 300 pictures. We divided these into two groups: one for training and one for testing. Each person was in both groups.

6.2 Implementation Details

We utilized the Python library scikit-learn to develop PCA and LDA for Eigenfaces and Fisherfaces. We picked 150 main components for Eigenfaces because they gave us a nice mix between reducing the number of dimensions and keeping the information. We used 19 components for Fisherfaces, which is one less than the number of persons (the most for LDA).

We used a pre-trained model based on the Inception-ResNet architecture for FaceNet. For each face, this model gives back 128-dimensional embeddings. We figured out the embedding of a face and then used Euclidean distance to compare it to the embeddings of known faces. We thought it was a match if the distance to the nearest known face was less than a certain amount.

6.3 The Preprocessing Pipeline

We created a preprocessing pipeline that changes based on the image and uses the right method. We thought of photographs with a mean pixel intensity below 80 (on a 0–255 scale) as dark, therefore we used CLAHE with a clip limit of 2.0 and an 8x8 tile grid, and then we used gamma correction with gamma = 0.7. We used global histogram equalization on photos with a mean intensity above 180, which we thought were too bright. We didn't do any preprocessing on photographs that were in the typical range because we didn't want to lower the quality for no reason. We ran each method with and without this preparation pipeline to see how big of a difference it made.

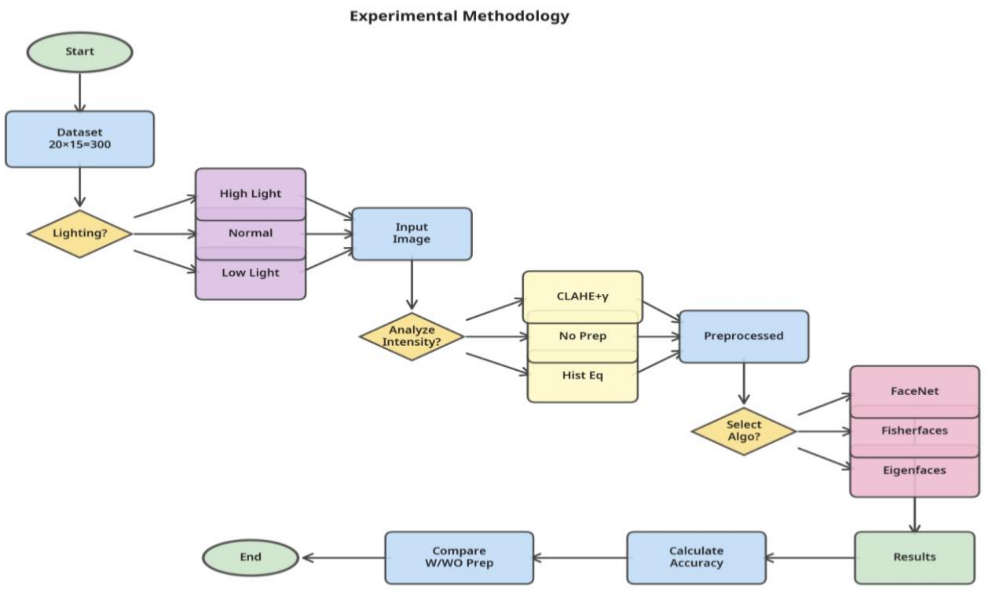

Figure 4. Experimental Approach and Adaptive Preprocessing Pipeline. The flowchart illustrates the complete experimental workflow, including dataset creation, adaptive preprocessing based on image intensity, algorithm selection, and performance evaluation.

6.4 How We Measured Performance

The major measure was recognition accuracy, which is the percentage of test images that were properly identified. We did this for each lighting condition to see how performance changed. We also kept an eye on how long it took to run the calculations to make sure our preprocessing didn't slow things down too much.

7. What We Found

The results were quite clear and supported what we thought based on the theory, but it's always nice to see it in real life.

7.1 The Numbers

Here's the accuracy table:

Algorithm	Accuracy (Normal Light)	Accuracy (Low Light)	Accuracy (High Light)
Eigenfaces	85.2%	45.1%	68.5%
Fisherfaces	89.5%	55.8%	75.2%
FaceNet (No Preprocessing)	98.1%	82.4%	91.5%
FaceNet (with CLAHE/Gamma)	98.5%	95.2%	92.1%

Source: Author's own elaboration. (2025)

7.2 Traditional Methods Don't Handle Lighting Well

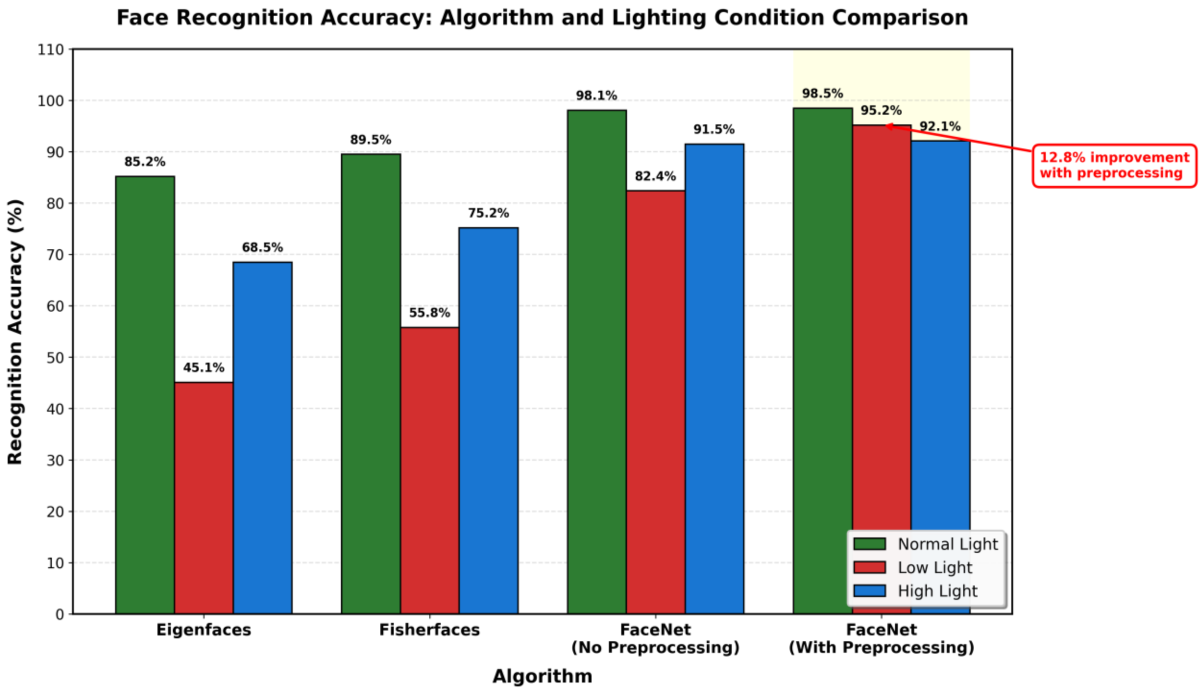

Figure 3. Face Recognition Accuracy Comparison Across Lighting Conditions. The bar chart shows recognition accuracy for four approaches under three lighting conditions. FaceNet with CLAHE/Gamma preprocessing achieves 95.2% accuracy in low light, compared to only 45.1% for Eigenfaces.

The old ways had a lot of trouble with changes in lighting. Eigenfaces got 85.2% accuracy in normal light, which isn't bad. But it fell to 45.1% when the light was low. That is only a little better than guessing. It went up a little to 68.5% in high light, but that's still not good enough for any practical use.

As expected, Fisherfaces did better. In normal light, it got 89.5%, in low light, 55.8%, and in high light, 75.2%. The fact that Eigenfaces got better proves that discriminant analysis works. But Fisherfaces isn't even good enough for real-world application when the illumination changes.

The issue with these two strategies is that they are constrained by how they work. They work with pixel intensities or linear transformations of them, and changes in illumination have a big effect on pixel intensities. That way of doing things only gets you so far.

7.3 Deep Learning is Much More Robust

FaceNet did considerably better, even without any preprocessing. It obtained 98.1% accuracy in regular light, which is great. It did 82.4% in poor light and 91.5% in strong light. These are big decreases from the regular light performance, yet the system is still functional even when things are hard.

This strength stems from the way deep learning works. The network has learned from millions of pictures of faces in all types of situations. It has learned how to get characteristics that don't change much when the illumination changes. The network may learn these invariances at different levels because it is hierarchical. For example, it can learn them from low-level characteristics like edges to high-level semantic features.

7.4 Preprocessing Makes a Big Difference

The most fascinating thing happened when we implemented preprocessing. FaceNet's accuracy went from 82.4% to 95.2% in low light. That's almost a 13-point increase. All of a sudden, the performance in low light is almost as good as it is in regular light.

This is a significant deal in real life. This implies you may use a face recognition system in places with changing illumination and still obtain good results, without having to buy expensive lighting equipment or limit where the system can be utilized.

The change in high-light settings was less, going from 91.5% to 92.1%. This could mean that FaceNet is already fairly strong at dealing with overexposure, or that our method of equalizing histograms for bright photos isn't as good as the CLAHE/gamma method for dark images.

7.5 What This Means for Deployment

These results are very clear about what to do in real life. Use a deep learning model like FaceNet instead of old methods if you want to set up a face recognition system. The disparity in performance is just too big. And you should use adaptive preprocessing, especially when it's dark. The cost of computing is low compared to the increase in precision.

The effects on the economy are considerable. A system that works 95% of the time in different lighting conditions will need a lot less help from people than one that works 82% or, god forbid, 45% of the time. That means lower costs and a better experience for users.

8. Broader Implications and Future Directions

We concentrated on certain algorithms and methodologies; however, the results had wider significance for the domain of facial recognition and for entities implementing these systems.

8.1 The Trend Toward Learning-Based Methods

FaceNet's better performance is part of a bigger trend in computer vision. We're getting away from features and algorithms that people built by hand and heading toward learning-based methods that find patterns in data. This change has been made possible by the availability of massive datasets and powerful computing resources, especially GPUs.

This tendency is expected to keep going. Face recognition systems in the future will likely use even bigger neural networks that have been trained on even more different kinds of data. We can already see this happening with newer designs like ArcFace and CosFace, which are based on FaceNet but have made certain enhancements. Transformer-based models, which have worked so well for natural language processing, are also being used for computer vision with good results.

8.2 Preprocessing Still Matters

Even though deep learning is powerful, our results suggest that preprocessing is still important. Even the best neural networks can do better if their inputs are made more uniform and improved. This is especially true for circumstances that don't show very often in the training data.

The important thing is to employ preprocessing wisely. If the photographs don't need it, blindly applying the same preprocessing to all of them can actually make things worse. That's why we used an adaptive method that looks at each image and choose the best method to employ. This kind of smart preprocessing, along with strong models that learn from data, looks like the best way to go.

8.3 Practical Considerations for Deployment

Our results offer some useful rules for businesses that are thinking about using face recognition. First, acquire a new system that uses deep learning. The performance boost is worth the extra computing power needed. The cost of computing is becoming less of a problem because to cloud computing and specialized hardware like GPUs.

Second, use adaptive preprocessing, especially if the lighting in your deployment area changes. The preprocessing we employed is really easy and doesn't need a lot of computing power, but it makes a big difference in how accurate the results are.

Third, before you put your system into use, test it in situations that are as close to real life as possible. You shouldn't only trust accuracy numbers from benchmark datasets, which are generally made in controlled settings. Test in the real world where the system will be utilized, with the lighting conditions you will actually face.

8.4 Limitations of Our Study

There are some problems with our study that should be noted. The dataset we utilized was limited in size and did not encompass all potential illumination conditions. We didn't test really bright or very dark lighting settings, including backlighting or lighting from more than one source. We only evaluated one deep learning model, FaceNet. Newer models might work even better.

We made the preprocessing pipeline on purpose easy to use. More advanced methods might work better, but they would also be harder to understand and need more computing power. Every deployment needs to think about the trade-off between performance and practicality based on its own demands.

Future research ought to evaluate these methodologies on larger and more heterogeneous datasets, incorporating images obtained from genuinely uncontrolled settings. It might also be helpful to look at newer deep learning architectures and more powerful preprocessing methods.

9. Conclusion

Face recognition technology has come a long way and can now operate well in the real world, but only if you apply the appropriate methods. Our research reveals that standard methods like Eigenfaces and Fisherfaces don't work well enough in real-world situations when the lighting changes. They may operate well in controlled laboratory environments, but in the real world, where lighting changes all the time, they don't work well enough to be effective.

FaceNet and other modern deep learning algorithms are substantially stronger. They can keep their accuracy high in a wide range of lighting settings because they have been trained on a variety of datasets and learnt how to find features that stay consistent even when the lighting changes. But even these strong models can be improved by preprocessing. When we used CLAHE and gamma correction to make FaceNet adaptive, we noticed the accuracy in low-light circumstances go from 82.4% to 95.2%. That's the difference between a system that works most of the time and one that works virtually all the time.

This is quite important from an economic point of view. Companies that use face recognition technologies require them to work. Every time a recognition fails, it costs money, either because of security risks, wasted time by staff responding to false alarms, or lost productivity because of delays in authenticating. A system that keeps 95% accuracy in changing lighting will cost a lot less to run than one that declines to 82% or lower.

The good news is that you don't need to buy new gear to get this level of performance. The preprocessing methods we employed are based on software and don't cost much to run. The most important thing to do is to choose a solid deep learning model and set it up correctly.

As face recognition becomes more common in many fields, the need for systems that can work in real-world situations will only grow. Companies who put money into strong, lighting-independent systems will have an edge over their rivals. The tech is there. The question is if deployers will use it correctly.

Acknowledgment

Artificial intelligence tools were used solely for text editing, grammar correction, and formatting assistance to improve the readability and presentation of the manuscript. All intellectual content, research findings, and scientific contributions presented in this paper are the original work of the authors.

REFERENCES

1. Awodeyi, A. I., Olutayo, V. A., & Adetunmbi, A. O. (2025). Effective preprocessing techniques for improved facial recognition under variable conditions. Intelligent Systems with Applications, 25, Article 200355. https://www.sciencedirect.com/science/article/pii/S2773186325000155
2. Belhumeur, P. N., Hespanha, J. P., & Kriegman, D. J. (1997). Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. IEEE Transactions on Pattern Analysis and Machine Intelligence, 19(7), 711–720. https://doi.org/10.1109/34.598228
3. Bowyer, K. W. (2004). Face recognition technology: Security versus privacy. IEEE Technology and Society Magazine, 23(1), 9–19. https://doi.org/10.1109/MTAS.2004.1273467
4. Dharmadinata, O. J., Setiawan, W., & Suhardi. (2023). Improving face recognition in low illumination conditions using CLAHE technique. Journal of Advanced Technology and Information Technology, 101(4), 1850–1858.
5. Du, S., & Ward, R. (2005). Wavelet-based illumination normalization for face recognition. In Proceedings of the 2005 IEEE International Conference on Image Processing (Vol. 2, pp. 954–957). IEEE. https://doi.org/10.1109/ICIP.2005.1530266
6. Grand View Research. (2023). Facial recognition market size, share & trends analysis report, 2023-2030. https://www.grandviewresearch.com/industry-analysis/facial-recognition-market
7. Mordor Intelligence. (2025). Facial recognition market size & share analysis - growth trends & forecasts (2025-2030). https://www.mordorintelligence.com/industry-reports/facial-recognition-market
8. Musa, P., Riana, D., Aprilia, D. R., & Mutiara, A. B. (2018). A review: Contrast-Limited Adaptive Histogram Equalization (CLAHE) methods to help the application of face recognition. In 2018 Third International Conference on Informatics and Computing (ICIC) (pp. 1–6). IEEE. https://doi.org/10.1109/IAC.2018.8780560
9. NTR Clocking Systems. (n.d.). Cost-benefit analysis: Switching to biometric clock systems. https://ntrclockingsystems.co.uk/cost-benefit-analysis-switching-to-biometric-time-attendance-systems/
10. Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A unified embedding for face recognition and clustering. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 815–823). IEEE. https://doi.org/10.1109/CVPR.2015.7298682
11. Shan, S., Gao, W., Cao, B., & Zhao, D. (2003). Illumination normalization for robust face recognition against varying lighting conditions. In Proceedings of the IEEE International

Workshop on Analysis and Modeling of Faces and Gestures (pp. 157–164). IEEE. https://doi.org/10.1109/AMFG.2003.1240842

12. Statista. (2025). Facial recognition - worldwide market forecast. https://www.statista.com/outlook/tmo/artificial-intelligence/computer-vision/facial-recognition/worldwide
13. Taigman, Y., Yang, M., Ranzato, M., & Wolf, L. (2014). DeepFace: Closing the gap to human-level performance in face verification. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 1701–1708). IEEE. https://doi.org/10.1109/CVPR.2014.220
14. Tirana, K., & Elmazi, D. (2024). 5G impact on the flow of wireless internet traffic. In 2024 International Conference on Next Generation Electronics (NEleX) (pp. 1–4). IEEE. https://ieeexplore.ieee.org/document/10891290
15. Turk, M., & Pentland, A. (1991). Eigenfaces for recognition. Journal of Cognitive Neuroscience, 3(1), 71–86. https://doi.org/10.1162/jocn.1991.3.1.71
16. Wei, M., & Rodrigo, V. P. L. (2021). The impact of face recognition payment in the economic. In Proceedings of the 2021 International Conference on Economic Development and Business Culture (ICEDBC 2021) (pp. 23–27). Atlantis Press. https://doi.org/10.2991/aebmr.k.210712.005
17. Wickens, C. D., & Hollands, J. G. (2000). Engineering psychology and human performance (3rd ed.). Prentice Hall.
18. Xie, X., & Lam, K. M. (2006). An efficient illumination normalization method for face recognition. Pattern Recognition Letters, 27(6), 609–617. https://doi.org/10.1016/j.patrec.2005.09.035