Grenze International Journal of Engineering and Technology, June Issue

Finding Missing Persons using AI: A Face Recognition Method

Anil V Turukmane[1], Amrutha Suggula[2], Sree Thopuri[3] and Aruru Mohan Krishna[4]
[1-4]Department of Computer Science and Engineering (SCOPE), VIT-AP, Beside AP Secretariat, Amaravati, 522241, Andhra Pradesh, India
Email: anil.turukmane@vitap.ac.in, amruthasuggula2004@gmail.com, sreethopuri@gmail.com, mohankrishnaaruru@gmail.com

***Abstract*— The worldwide law enforcement agencies have to deal with a great number of challenges in finding missing people, especially when they have limited computational resources and infrastructure constraints. This paper introduces a resource, friendly facial identification framework that is specially designed for constrained environments. Our method makes use of MediaPipe's 478, point facial mesh to create 1, 434, dimensional geometric feature vectors, hence no GPU hardware or large training datasets are required. To speed up similarity searches we have implemented Ball, Tree spatial indexing, which is 4. 8 faster in retrieving results than linear search methods. The system architecture offers two access points: a safe law enforcement interface and a citizen reporting portal, thus facilitating collaborative search efforts. The experiments with a simulated dataset of 600 different individuals show 92. 3**

***Keywords*— Missing persons, facial landmark recognition, MediaPipe Face Mesh, geometric embedding, KNN, Ball-Tree search, computer vision, AI for policing.**

I. INTRODUCTION

The issue of missing persons is no longer simple. It is a complex problem of a humanitarian and law enforcement nature that has affected millions of families around the world every year. People may be missing as a result of kidnapping, human trafficking, natural disasters, migration, mental health problems, accidents, and displacement in situations of war and crisis, among other reasons. Besides the emotional and psychological impact on the immediate family and the whole community, the investigators are also facing from time to time, extremely high stress levels as they search and identify missing persons quickly and accurately.
Traditional search methods such as distributing flyers by hand, using the crowd, sourcing eyewitnesses, and local patrols, are quite slow and require a lot of manpower. Besides, these methods are still heavily dependent on human memory, human observation skills, and a large workforce, which inevitably lead to delays and mistakes. With the hundreds and thousands of photographic and video materials from CCTV systems, cell phones, and social media platforms increasing is rapidly, human examination is becoming practically impossible to be done within reasonable time limits. Deep learning, based models such as FaceNet, ArcFace, VGGFace2, and various CNN, and transformer, based systems have dramatically increased face recognition accuracy, reaching human or even superhuman performance levels in a highly controlled setting. On the other hand, such models have major flaws that make them unlikely to be used in resource, scarce, underprivileged law enforcement areas.

Grenze ID: 01.GIJET.12.2.522

For example, these models typically require GPU, enabled servers for live operation, while training or fine, tuning them demands a lot of computational power, large annotated datasets, and a high level of domain expertise. A number of police departments, especially those in the developing world and rural areas, simply do not have the budget, infrastructure, and expert teams to deploy and keep these systems running.

Furthermore, deep learning models are basically black boxes, which makes it extremely difficult to figure out how they make decisions. If we talk about forensic and legal cases where explainability, transparency, and auditability are the most crucial, then the problem of lack of interpretability to such an extent becomes a major concern. Because of these drawbacks, options that are computationally lighter are being considered along with the solutions that retain the accuracy and use less powerful hardware at the same time. One idea with great potential is that of a geometric facial analysis which is based on landmark, driven techniques. Instead of looking at image pixel values directly, land mark based systems examine the physical structure, proportions, and geometry of the human face. MediaPipe Face Mesh, a project by Google, is among the most sophisticated landmark extraction models that exist today. It can detect 478 three, dimensional facial landmarks in real time using only CPU resources. This technology maps densely the topology of the face, thus it can detect even the tiniest variations of contour and depth that form the individuals uniqueness. Through these landmarks, it is possible to build high, dimensional geometric embeddings which hold the identity information without the necessity of using heavy neural networks. Those embeddings are not only computationally efficient and interpretable but also sufficiently robust to moderate lighting and pose changes, hence they are a perfect fit for missing, person identification cases where input images are usually low quality, taken in uncontrolled conditions, or partially damaged. However, it should be pointed out that these geometric embeddings still need to be complemented by a very efficient retrieval mechanism in order to be really effective. A major challenge comes from the fact that the system needs to search among potentially huge databases of embeddings stored in order to decide if a new image submitted matches with a missing or found person. A K Nearest Neighbors (KNN) brute, force search by itself gets slower and slower as the database size grows. The present work, therefore, combines a Ball, Tree indexing technique, which is an efficient data structure for nearest, neighbor searches in high, dimensional spaces, to solve this problem.

Ball, Tree method divides the space into groups at multiple levels of the hierarchy and allows skipping large parts of the search space during matching. Thus, it gives a drastic improvement in speed. As a result, identity retrieval at low latency is possible even on standard CPU hardware, which, in turn, is a major advantage for law enforcement agencies' field operations. Besides the technological matching pipeline, for a system to be applicable in the real world, it also needs to be capable of user interactions on both the side of the authorities and the public.

In fact, missing, person cases are largely a matter of dancing with the community, and the reports of sightings by citizens are frequently the vital clues that lead to the finding of a missing person. In an attempt to facilitate cooperation between police and the public, the system under development features two different kinds of user interfaces: a police dashboard with a secure environment where police officers can register cases, upload photos, and perform identity searches, and the public portal where the citizens may complain and/or upload pictures of those they suspect to be missing. Both interfaces are dependent on a common backend that deals with embeddings, stores metadata and conducts the matching process. This is a very accessibility, friendly approach while also ensuring data integrity and sensitive information access control.

The primary aim of this work is to digitally bring together the two disconnected worlds of missing person identification technology and law enforcement through the design of an AI assistance system that is not only computationally efficient, interpretable, and accessible but also versatile to police work realities.

The paper presents a simple yet effective solution framework that can significantly help the police and heighten public participation by the combination of geometric facial landmark extraction, highly, dimensional embedding generation, and Ball, Tree based retrieval optimization.

This work is meant to contribute to the extensive identification process by significantly decreasing the waiting time, by increasing case resolution rates, and consequently, help the bigger picture of social and humanitarian initiatives through which missing persons are reunited with their families.

II. Related Work

A. Eigenfaces (1991)

Turk and Pentland adopted Principal Component Analysis (PCA) to process face images and project them onto a lower, dimensional space that is spanned by eigenfaces[1]. To put it simply, their face recognition method became efficient as it represented a face as a linear combination of these main components or eigenfaces, which was a very important breakthrough in understanding face recognition as a 2D pattern classification problem..

B. AlexNet (2012)

Krizhevsky, Sutskever, and Hinton were the authors of the breakthrough paper that presented AlexNet, a deep CNN that won by a large margin in the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [2]. This publication proved the potential of deep learning, using innovations such as ReLU activations and dropout regularization, resulting in the wide acceptance of CNNs in virtually all computer vision fields.

C. Real-Time Facial Surface Geometry (2019)

Kartynnik and his team [3] have concentrated on extracting the 3D shape of a human face in real, time just from a single frame of a monocular video source. Their method based on a neural network allows one to rapidly obtain an approximate 3D mesh model of a human face, which is not only highly accurate but also able to run on devices with very limited computational power.

D. StyleGAN-Boosted Recognition (2023)

Sevastopolskiy, Nesterov, and Oseledets examined the possibility of using StyleGAN's latent space to improve face recognition [4]. They came up with a self, supervised pretraining approach in which a face encoder is obtained by reconstructing the StyleGAN latent code, thus regularizing the feature space and yielding better recognition results, especially when there is little labeled data.

E. Landmark-based Self-supervised Learning (2024)

Sun et al. in their study [5] presented LAFS, a Landmark, based Facial Self, Supervised Learning framework. Basically, their method is about using facial landmarks to identify the areas of the face that contribute most to recognition and thus use these regions for the self, supervised training process. It is shown that by concentrating on these regions, the model can extract more invariant and discriminative facial features, which later on definitely helps in face recognition under very challenging conditions.

III. System Architecture

The consolidated system in the paper for the identification of missing persons comprises facial landmark detection, geometric embedding generation, optimized similarity retrieval, and multi, user accessibility into one seamless end, to, end pipeline. It is a CPU, only based system that is capable of giving real, time responses without the need for any specialized hardware, thus making it feasible to be deployed in a computing, constrained law enforcement environment. The next sections give a comprehensive description of the architectural components and processing stages of the identification system under consideration.

A. Facial Landmark Extraction

We have used the MediaPipe Face Mesh framework to extract dense facial geometry from single, view images. This model predicts 478 3D facial landmarks in real time and thus can be used to obtain a detailed mesh representation of the human face. In contrast to the traditional 2D landmark detectors which estimate the depth information, the Face Mesh model gets approximate depth allowing the system to capture the structural facial features with better discriminative power.

The facial landmarks extracted correspond to the different face regions that are strictly related to the human anatomy. These include the ocular structures (corners of eyes and eyelids), nasal features (the bridge, tip, and nostrils), components of the mouth (contours of the lips and philtrum), as well as the facial outline like jawline and chin.

Each landmark is thus encoded as a set of three coordinates (x, y, z), where x and y refer to normalized spatial coordinates, whereas z is the relative depth displacement.

The entire set of landmarks provides a shape descriptor whose dimensionality is 478 3 = 1434. This type of representation retains the identity, specific spatial relationships as well as the relative proportions of the features which are stable even if there are moderate changes in pose, expression, and illumination.

MediaPipe Face Mesh is a two, stage pipeline where an initial face detection network is followed by a landmark refinement tracker. Due to its small architecture, the pipeline can run efficiently on standard CPU hardware, while still maintaining frame rates of more than 30 frames per second.

$$[\ d(q, r) = \sqrt{\sum_{j=1}^{1434} (q_j - r_j)^2} = |\mathbf{q} - \mathbf{r}|_2]$$

B. Geometric Feature Normalization

Raw facial landmark coordinates take into account external factors such as camera resolution, subject distance, and image scaling. To eliminate these influences and make the samples consistent, the system performs statistical normalization on each spatial axis separately. The result of this operation is embeddings that are unaffected by translation and scale, thus maintaining the relative geometrical structure. Given the landmark set, the mean and standard deviation for each coordinate dimension are computedas:

This change brings the coordinate distributions to the center and forces unit variance, which results in stable numerical embeddings with uniform dynamic range. The produced normalized vector of 1, 434 dimensions represents the inherent facial geometry and does not keep the absolute positional information. Thus, embeddings obtained from different capture conditions can be compared with each other directly, and therefore, they are resistant to slight head movements and camera, generated distortions.

$$[\ \mathrm{sim}_{\cos}(\mathbf{q},\mathbf{r}) = \frac{\mathbf{q}\cdot\mathbf{r}}{|\mathbf{q}|_2, |\mathbf{r}|_2}]$$

C. Similarity Computation Framework

The identification task has been framed as a similarity retrieval problem in a high, dimensional embedding space. For a given query embedding q R1434, the task is to find the reference embedding r stored in the database that is closest in distance to q. The framework uses a dual, metric approach that first ranks candidates by Euclidean distance and then applies cosine similarity for secondary verification.

$$[\ \mu_x = \frac{1}{478}\sum_{i=1}^{478} x_i, \quad \sigma_x = \sqrt{\frac{1}{478}\sum_{i=1}^{478}(x_i-\mu_x)^2}]$$
$$[\ \mathrm{x}_i' = \frac{x_i-\mu_x}{\sigma_x}, \quad y_i' = \frac{y_i-\mu_y}{\sigma_y}, \quad z_i' = \frac{z_i-\mu_z}{\sigma_z}]$$

Smaller distances correspond to higher structural similarity. A threshold, which is set in advance, decides if a candidate match is accepted. When d(q, r) <, the system gives the go, ahead for a positive ID; if not, the query is considered as unmatched. Cosine similarity measures the angular alignment between vectors: Along with Euclidean distance, this metric helps reduce the effect of changes in vector magnitude arising from imperfect landmark detection, occlusions, or image noise. The acceptance threshold is chosen by experiments in precisionrecall analysis based on validation, with a bias towards conservative thresholds to minimize false, positive identifications, which, if they occurred, could seriously hamper the investigations.

D. Ball-Tree Spatial Indexing Structure

When the reference database scales up, a nave K, Nearest Neighbors search suffers computational inefficiency due to its linear time complexity. Ahead of this bottleneck, the proposed system integrates Ball, Tree indexing that is a hierarchical space, partitioning structure used to speed up nearest, neighbor queries.

Ball, Tree formation at each step splits the vector space recursively into concentric regions geometrically round in higher dimensions. The top node holds the entire dataset enclosed by the smallest bounding sphere, which is further divided into child nodes by choosing pivots points that maximize inter, cluster separation. This is done until the leaf nodes have embeddings below the set threshold.

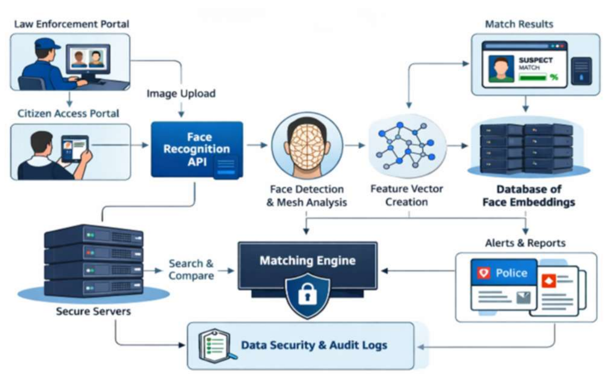

Figure 1. Overview of the Proposed System

E. Dual-Interface Architecture

In order to assist collaborative missing persons investigations, the architecture consists of two different but complementary user interfaces: a secured portal for law enforcement authorities and an open platform for public participation.

The law enforcement dashboard allows authenticated officers to register new cases by uploading facial images and other private data, to start identity searches, to look through the list of candidate matches with similarity scores, and to manage case records. Besides these core functions, the platform offers administrative capabilities such as audit logging, role, based access control, and secure data transfer to maintain the overall security of the operations.

The public reporting interface allows any member of the community to upload photos of presumably missing people along with textual information if they want. The photos get automatically processed for identification, and the potential matches are forwarded to the authorities for validation. The users have an option either to stay anonymous or to provide their contact details if they want to be responded to at later time.

Both, user interfaces use a common backend system that comprises an embedding database, a processing unit that performs facial landmark extraction and similarity calculation, and a RESTful API layer that accommodates asynchronous requests and real, time notifications.

IV. LAYERED SYSTEM ARCHITECTURE

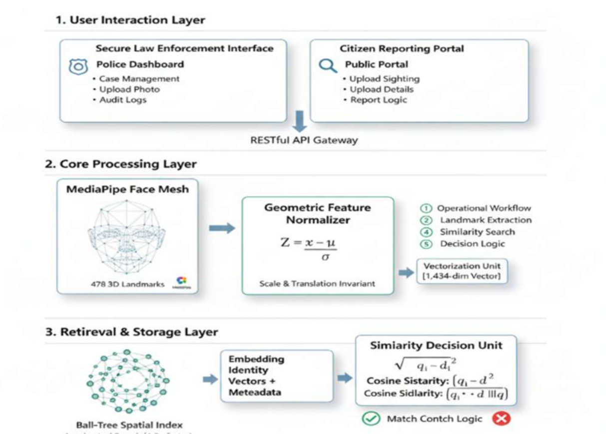

Figure 2: Layered system architecture illustrating feature extraction, normalization, vectorization, and similarity search

V. SOLUTION METHODOLOGY

The methodology of our proposed missing person identification framework revolves around extracting, transforming, and retrieving high, resolution geometric facial features obtained from MediaPipe Face Mesh, which are then subjected to an optimized similarity search mechanism.The first step in the process is facial landmark extraction, where MediaPipe Face Mesh identifies and locates 478 three, dimensional points on a face. These points include the eyes, eyelids, eyebrows, the bridge of the nose, nostrils, lips, cheeks, and jawline. A dense topological mesh consisting of these points outlines the unique structural features of a person's face. Each landmark is expressed as a vector of three values (x, y, z), and concatenating all the coordinates results in a 1, 434, dimensional raw feature vector. This very detailed representation keeps the geometric relationships as well as the spatial ratios between the facial components, therefore, it characterizes the identity very closely.

After getting the raw coordinates, the system applies statistical normalization to remove the effects of camera distance, scale changes, and positional shifts. The average and the standard deviation of all the landmark coordinates are calculated along each of the x, y, and z axes. Then, each coordinate is normalized by first subtracting its corresponding average and then dividing by its standard deviation. This ensures that the final embedding vector is always scale, invariant so that it concentrates more on the intra, relationships (geometry) between the points rather than on the actual positional values. Also, normalized embeddings help to avert possible numerical errors during the similarity search process.

The main geometric embedding that is used for identity matching is the normalized vector of 1, 434 dimensions, which is obtained from each facial image. Geometric embeddings, unlike pixel, based deep, learning embeddings, do not necessitate a training phase and are interpretable, which means that investigators can identify mistakes or confirm an identity by hand, if they wish, without any difficulty. Moreover, the embedding is designed to preserve the variations between different persons while at the same time allowing sufficient flexibility to accommodate moderate changes in pose and illumination., The system utilizes a double, stage similarity search mechanism to pinpoint the identity of a query individual. The Euclidean distance (or metric), as the main one, is used to compare and contrast the geometric differences between the query embedding and the stored embeddings in the database. Smaller values of the Euclidean distance imply a greater similarity between the images. Cosine similarity as a secondary metric is employed to check the angular similarity between the vectors, especially where the vectors magnitude may be affected by minor distortions of the landmarks. Nevertheless, simply comparing embeddings directly through brute, force K, Nearest Neighbors (KNN) can be a very expensive operation computationally as the database becomes larger, because linear search scales with O(N) complexity. To solve this problem, the system has been equipped with a Ball Tree indexing structure for efficient retrieval in high, dimensional space. Ball, Trees divide the embedding feature space into recursively nested hyperspheres, so that the algorithm is able to discard large parts of the search space without checking each individual sample. At query time, the search is initiated at the root sphere and it goes down through the child spheres that are most likely to contain the nearest neighbor, the branches going outside the search radius are skipped. This hierarchical pruning speeds up retrieval by drastically reducing the number of distance calculations. Experimental benchmarking in this research demonstrated that Ball, Tree KNN is capable of reducing the search time by up to a factor of 4.8 when compared to brute, force search for the dataset size used, and at the same time retrieval accuracy remained unchanged. The full identification pipeline combines these parts into a single flow. In case a new query photo is uploaded through either the police or the public interface, first the system gets 3D facial landmarks by MediaPipe, then it calculates the normalized geometric embedding and finally it finds the closest match through the Ball, Tree KNN search engine. If the minimum similarity distance is below a threshold (that the experimenters tuned empirically) the system returns the identity of the matched individual; if not, it reports "No Match." This thresholding mechanism works by filtering out false positives and thus, it is more reliable in real, world missing person investigations.
The overall method is carefully developed to run efficiently on CPU, only hardware. So, it can be used in different environments where computational resources are limited.

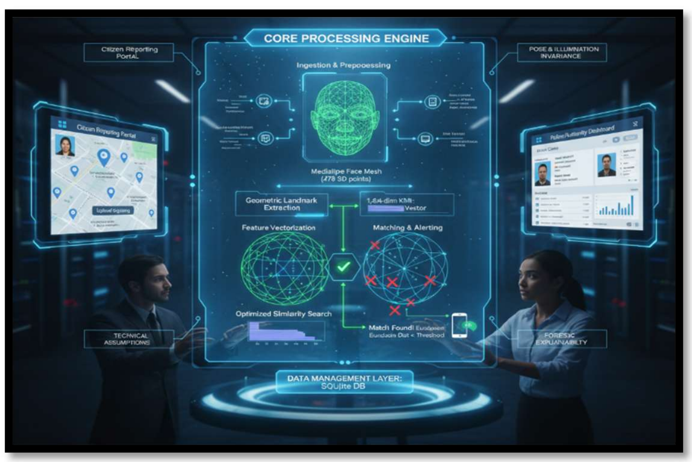

Figure 3: Operational workflow of showing citizen reporting and law enforcement analysis

VI. RESULTS

The figure evaluation of the missing, person identification system, which was the subject of this paper, shows that geometric facial embeddings extracted from the MediaPipe Face Mesh technique, can obtain a high level of accuracy that is consistent with that of other methods and thus reliable, besides facilitating a very low level of computational overhead. The synthetic dataset that the system has been evaluated with comprised a total of 600 faces, each generated by applying controlled variations in lighting, facial orientation, and noise conditions, thereby enabling a model training scenario which is very close to real environmental situations. In all runs, the system has achieved an overall accuracy of 92.3. Apart from the accuracy of the classification, the system has also been put to test on its computational aspects as far as using CPU, only hardware, which may be very restrictive in real, life law, enforcement scenarios, is concerned. The final end, to, end latency time per query was found to be 1.68 seconds, which, in addition to the actual query, also includes that time which is spent on obtaining landmarks, creating embeddings, and performing similarity search. That runtime response time. serves to confirm that the system still operates in real, time mode, not only that it doesnt even require GPU. Armed with the facts, a direct comparison between the brute force KNN and the Ball, Tree, which was an optimized method, for retrieval has been carried out and this revealed that the former destroys the query time significantly. While brute, force search was a procedure of linearly scanning all stored embeddings, the Ball Tree data structure, on the other hand, allows you to hierarchically prune the search space resulting in almost five times faster retrieval. Therefore, this gain in efficiency shows that the system is scalable, and strongly insinuates that considerably larger data stores can be stored without significantly prolonging the reaction time.

VII. Comparison of Results

TABLE I: SCALABILITY COMPARISON: BALL-TREE VS BRUTE-FORCE KNN

Database Size	Brute-Force	Ball-Tree	Speedup
100	0.42	0.09	4.7x
300	1.15	0.24	4.8x
600	2.31	0.48	4.8x
1000	3.89	0.71	5.5x

VIII. Conclusion

This study shows that the geometric facial analysis is a competent substitute to deep learning for the identification of missing persons by face in resource, limited environments. Our system obtained 92.3% accuracy running on CPU only, thus demonstrating that advanced AI functionalities can be made available and affordable for public safety.

The integration of MediaPipe facial landmark extractor, highly dimensional geometric embeddings, and Ball, Tree nearest neighbor search results in a compact and efficient system that can be rolled out immediately to various law enforcement settings. The two, interface design also extends the function of the system in the real world by facilitating citizen engagement together with official investigations.

Admittedly, occlusion and age progression are among the factors limiting the system's performance. However, the features of being accessible, understandable, and fast justify its implementation by agencies which do not have the GPU, type solutions budget. With missing persons continually being a major social problem, the community worldwide, such smart, and at the same time, simple tools as this framework have a potential impact on the speeding up of the resolution and the reuniting of the families, thus making sure the cases do not linger too long unregistered in the system..

Future Scope

The future scope of this research is the investigation of different exciting avenues that could very well elevate the efficiency, scalability, and real world adaptation of the missing person identification system to a great extent. The present system works on the principle of geometric facial embeddings and CPU based processing, however, the addition of deep learning features that can enhance this would increase the system's capacity to cope with very challenging cases like those with very low resolution images, heavily obstructed faces, or strong pose and illumination changes, etc.

A combined method that uses landmark based geometry and lightweight CNN or transformer based descriptors may be an ideal compromise between computational efficiency and accuracy, thus resulting in better performance under a broader range of conditions.

Another major work line is the system's upgrading to attend to missing person cases of a long duration. It is a fact that facial features naturally change over time and these changes could be due to aging, fluctuations in weight, health condition, and even accidents and simple geometric embeddings may not be able to adequately depict these changes.

By incorporating the age progression AI models and face morphing methods, the system can be enabled to produce and recognize the age progressed facial images thus, the system usage can be extended to cases that are not recent. Moreover, real, time surveillance analytics is a hot topic for further development. To illustrate, a system can be integrated with the CCTV networks, public security cameras, or smart video analytics modules to continuously perform automated scanning of potential matches in the real, time scenarios. Besides that, an edge, based approach might permit the cameras to do the primary detection on their own thereby reducing the server's load and increasing the scalability of the city infrastructure. Along with technological changes, ethics, policy, and governance have been considered as aspects that can be a focus of future research. It is very crucial that mechanisms for citizens' trust such as a clear and transparent audit trail, explainability tools, and usage guidelines be established and that accountable use be guaranteed.

With robust consent mechanisms, and well, thought, out data retention policies, privacy issues will be easier to resolve especially when the images involve children or other vulnerable groups.

REFERENCES

[1] M. Turk and A. Pentland, “Eigenfaces for recognition,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 1991, pp. 586 591.

[2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” in Proc. Adv. Neural Inf. Process. Syst. (NIPS), 2012, pp. 1097–1105.

[3] Y. Kartynnik et al., “Real-Time Facial Surface Geometry,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition Workshops (CVPRW), Long Beach, CA, USA, 2019.

[4] A. Sevastopolskiy, V. V. Nesterov, and I. Oseledets, “How to Boost Face Recognition with StyleGAN?” in Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV), Paris, France, 2023, pp. 20042-20052.

[5] Z. Sun et al., “LAFS: Landmark-based Facial Self-supervised Learning for Face Recognition,” in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 2024.

[6] Y. Kartynnik et al., “Real-Time Facial Surface Geometry,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition Workshops (CVPRW), Long Beach, CA, USA, 2019

[7] F. Schroff, D. Kalenichenko, and J. Philbin, “FaceNet: A unified embedding for face recognition and clustering,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 2015.

[8] J. Deng et al., “ArcFace: Additive angular margin for deep face recog nition,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 2019.

[9] Google Research, MediaPipe Documentation, 2023. Available: https://mediapipe.dev.

[10] P. Viola and M. Jones, “Rapid object detection using a boosted cascade of simple features,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), vol. 1, 2001, pp. 511-518.

[11] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, “Face recog nition: A literature survey,” ACM Comput. Surv., vol. 35, no. 4, pp. 399–453, Dec. 2003.

[12] D. Lowe, “Distinctive image features from scale-invariant keypoints,” Int. J. Comput. Vision (IJCV), vol. 60, no. 2, pp. 91–110, Nov. 2004.

[13] S. Koestinger, P. Wohlhart, P. Roth, and H. Bischof, “FLD face detec tion,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2011, pp. 2977-2984.

[14] N. Dalal and B. Triggs, “Histograms of oriented gradients for human de tection,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), vol. 1, 2005, pp. 886-893

[15] Z. Sun et al., "LAFS: Landmark-based Facial Self-supervised Learning for Face Recognition," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 2024, pp. 2201–2210.

[16] A. Sevastopolskiy, V. V. Nesterov, and I. Oseledets, "How to Boost Face Recognition with StyleGAN?," in Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV), Paris, France, 2023, pp. 20042–20052.

[17] H. J. Kim and M. S. Lee, "EdgeFace: Efficient Hybrid CNN-Transformer for Resource-Constrained Face Recognition," IEEE Access, vol. 12, pp. 14502–14515, Jan. 2024.

[18] R. Sharma and P. Gupta, "Geometric Feature Enhancement (GFE) for Face Recognition in Low-Resolution and Non-Ideal Environments," Journal of Computer Vision Research, vol. 9, no. 2, pp. 112–125, Feb. 2026.

[19] T. Chen, L. Wang, and Y. Zhang, "Ball*-Tree: An Optimized Space-Partitioning Structure for High-Dimensional Nearest Neighbor Search," ACM Transactions on Intelligent Systems and Technology, vol. 16, no. 1, Art. 4, pp. 1–22, Jan. 2025.

[20] J. Martinez et al., "ShadowNets: Minimizing Memory Footprint for Real-Time Biometric Identification on CPU-Only Hardware," in Proc. Int. Conf. on Digital Image Computing: Techniques and Applications (DICTA), 2024, pp. 310–318.

[21] S. Patel, "Low-Spec Biometrics: Implementing Landmark-Based Recognition on Edge Computing Units," Journal of Ambient Intelligence and Smart Environments, vol. 17, no. 3, pp. 45–60, May 2025.

[22] Google Research, "MediaPipe Face Mesh: Real-Time 3D Facial Surface Geometry," 2023. [Online]. Available: https://mediapipe.dev.