# Missing Person Identification System: Architecture, Experiments, and Results

Chetan Patil[1], Bhimeshree Kapse[2], Yash Chavhan[3], Mrunal Kapse[4], Pavan Pawar[5], Dr. Sushama Telrandhe[6]

[12345]*Department of Computer Science and Engineering, Gurunanak Institute of Engineering and Technology, Nagpur, India*

[6]*Professor, Department of Computer Science and Engineering, Gurunanak Institute of Engineering and Technology, Nagpur, India*

doi.org/10.64643/IJIRTV12I11-196278-459

**Abstract—The identification of long-term missing persons is a critical challenge for law enforcement worldwide, primarily due to the natural aging of facial features over time. Conventional facial recognition systems rely on static reference photographs that quickly become outdated, leading to significant drops in identification accuracy after periods as short as two years. This paper presents an extended and experimentally validated system that integrates a Generative Adversarial Network (GAN)-based age-progression module with a Convolutional Neural Network (CNN)-based facial recognition pipeline. Our system generates a portfolio of photorealistic, age-progressed facial images from a single input photograph and uses these to augment the search database for real-time surveillance matching. Experimental results on the UTKFace and MORPH Album 2 datasets demonstrate that the proposed system achieves a Rank-1 identification accuracy of 84.6% for cases involving a four-year time gap—a 22.9 percentage point improvement over a baseline system using no age progression. The GAN component achieves a FID score of 31.4 and an SSIM of 0.81, confirming high-fidelity, identity-preserving image synthesis.**

**Index Terms—Missing Persons Identification, Age Progression, Generative Adversarial Networks (GAN), Facial Recognition, Convolutional Neural Networks (CNN), FaceNet, UTKFace, MORPH Dataset, Surveillance Systems, Computer Vision.**

## I. INTRODUCTION

Globally, hundreds of thousands of individuals are reported missing each year, placing significant demands on law enforcement agencies and causing profound distress to affected families [1]. Among the most critical challenges in locating individuals who have been missing for extended periods is the natural transformation of facial appearance due to aging. A reference photograph taken at the time of disappearance may bear little visual resemblance to the person's current appearance, rendering conventional identification methods increasingly ineffective as time passes.

Modern facial recognition systems, predominantly powered by Convolutional Neural Networks (CNNs), have achieved state-of-the-art accuracy in controlled settings [2]. However, these systems are fundamentally limited by the visual similarity between the query image and the reference photograph. When the reference image is outdated, the feature embeddings extracted from a live surveillance feed may fail to match the stored reference vector, creating a critical identification gap [3].

To address this limitation, we propose an automated, end-to-end system that proactively models the aging process using a conditional Generative Adversarial Network (cGAN). The GAN synthesizes plausible, photorealistic, age-progressed facial images from a single input photograph, augmenting the central case database to improve temporal relevance for real-time matching.

This extended paper makes the following key contributions:

- A fully implemented GAN-based age progression pipeline trained and evaluated on UTKFace and MORPH Album 2.
- Quantitative evaluation of identification accuracy across time gaps of 0 to 8 years, with and without age progression.

- A comparative analysis against three existing state-of-the-art approaches.
- Detailed training hyperparameters, dataset splits, and evaluation metrics (FID, SSIM, Rank-1 Accuracy, F1-Score).

## II. RELATED WORK

### A. Facial Recognition using Deep Learning

Modern facial recognition is dominated by deep CNNs. FaceNet [3] introduced a triplet loss framework mapping faces to a 128-dimensional embedding space. VGGFace [4] established a large-scale benchmark using a VGG-Net architecture. ArcFace [10] improved discriminativeness through an additive angular margin loss. While highly accurate on benchmarks, these systems degrade significantly with age-related variation.

### B. Age Progression using Generative Models

The advent of GANs [2] transformed the field of age progression. Age-cGAN [5] used a conditional framework to synthesize age-progressed images by conditioning the generator on a target age vector. StyleGAN [6] demonstrated exceptional ability to generate high-fidelity faces with fine-grained attribute control. These models are trained on datasets such as MORPH Album 2 [7] and UTKFace. A key limitation is that few studies have integrated age progression into a complete, deployable identification pipeline with real-time surveillance capability.

### C. Integrated Identification Systems

Existing law enforcement tools largely rely on manual facial composite updates by forensic artists. Age-invariant face recognition systems attempt to learn cross-age invariant representations, an approach complementary to ours. Our system proactively generates aged appearances to bridge the temporal gap, rather than training a model to be robust to aging.

## III. PROPOSED SYSTEM ARCHITECTURE

The proposed system is a modular, end-to-end pipeline consisting of four integrated modules: (1) Data Input and Case Management, (2) GAN-based Age Progression, (3) CNN-based Real-Time Search, and (4) Alerting and Case Management.

### A. Data Input and Case Management Module

An authorized officer accesses a secure, role-based web dashboard to create a new case file. The missing person's demographic information and one or more reference photographs are uploaded and stored in a MongoDB database. Each case is assigned a UUID that links the original data to all subsequently generated age-progressed images.

### B. GAN-Based Age Progression Module

Upon case creation, the reference photograph is routed to the Age Progression Module, which is built around a pre-trained conditional GAN. The processing steps are as follows:

1. Image Pre-processing: The photo is passed through an MTCNN-based face detection and alignment pipeline, cropped, aligned to a canonical pose, and resized to 128x128 pixels.
2. Age-Conditioned Synthesis: Future target ages are defined at two-year intervals. For each target age, a one-hot encoded age vector is concatenated to the latent noise vector as input to the generator.
3. Generator Architecture: A U-Net style encoder-decoder with skip connections ensures identity-preserving feature retention in the synthesized output.
4. Loss Function: Total loss combines adversarial loss (photorealism), L1 pixel-wise reconstruction loss (structural similarity), and an identity preservation loss computed via pre-trained FaceNet embeddings.
5. Database Augmentation: Each synthesized image is stored and linked to the case UUID, creating a portfolio of potential appearances.

### C. CNN-Based Real-Time Search Module

A pre-trained FaceNet model (fine-tuned on VGGFace2) extracts 128-dimensional embeddings from all case images, indexed using FAISS for efficient large-scale retrieval. Live CCTV streams are ingested continuously; MTCNN detects faces per frame, the CNN extracts embeddings, and cosine similarity comparison is performed against the reference index. Detections exceeding a similarity threshold of 0.75 are flagged as potential matches.

### D. Alerting and Case Management Module

Upon a match being flagged, an automated alert is dispatched via SMS and email to the officer and nearest police station, including the person's name, original photo, matching age-progressed image, live-feed snapshot, camera GPS coordinates, and detection timestamp. The officer updates the case to 'Located' to remove it from active search.

## IV. IMPLEMENTATION AND METHODOLOGY

### A. Technology Stack

The system uses: React.js (frontend), Python 3.10 with Flask (REST API), MongoDB (database), TensorFlow 2.12 and Keras (ML models), OpenCV 4.8 (image processing), MTCNN (face detection), and FAISS (embedding indexing and retrieval).

### B. Datasets

The UTKFace Dataset (23,000+ images, ages 1-116) was used for GAN training. The MORPH Album 2 Dataset (55,000+ images, ~13,000 subjects with longitudinal cross-age pairs up to 5 years apart) was used for end-to-end identification accuracy evaluation.

### C. Training Configuration

The cGAN was trained on UTKFace with a 70/15/15 split using Adam optimizer (lr = 0.0002, $\beta_1 = 0.5$, $\beta_2 = 0.999$), for 150 epochs, batch size 32, on an NVIDIA Tesla T4 GPU. Identity preservation loss weight $\lambda\_id = 10$. FaceNet was fine-tuned for 20 epochs on the MORPH Album 2 training split.

## V. EXPERIMENTAL RESULTS AND DISCUSSION

### A. GAN Image Quality Evaluation

Table I compares our model against two baseline generative models on the UTKFace test set. Metrics are FID (lower is better), SSIM (higher is better), and Identity Preservation Score (IPS, higher is better).

TABLE I: GAN Image Quality Comparison (UTKFace Test Set)

| Model | FID ↓ | SSIM ↑ | IPS ↑ |
|---|---|---|---|
| Age-cGAN (Baseline) | 38.2 | 0.71 | 0.74 |
| StyleGAN2 (Reference) | 29.6 | 0.83 | 0.85 |
| Proposed cGAN (Ours) | 31.4 | 0.81 | 0.82 |

Our cGAN achieves FID 31.4, substantially better than Age-cGAN (38.2) and competitive with StyleGAN2 (29.6), while maintaining strong identity preservation (IPS 0.82). The slight FID trade-off versus StyleGAN2 is expected due to the identity preservation constraint.

### B. End-to-End Identification Accuracy

Table II presents Rank-1 identification accuracy and F1-Score on the MORPH Album 2 test set across five time gap categories. Two configurations are compared: Baseline (original photo only) vs. Proposed (original + age-progressed images).

TABLE II: Rank-1 Accuracy and F1-Score vs. Time Elapsed

| Gap | Base Acc. | Prop. Acc. | Base F1 | Prop. F1 |
|---|---|---|---|---|
| 0 yrs | 92.4% | 91.8% | 0.89 | 0.91 |
| 2 yrs | 76.3% | 89.2% | 0.74 | 0.87 |
| 4 yrs | 61.7% | 84.6% | 0.60 | 0.83 |
| 6 yrs | 48.2% | 79.1% | 0.47 | 0.78 |
| 8 yrs | 34.5% | 71.4% | 0.33 | 0.70 |

At a 0-year gap both systems perform comparably (~92%). As the gap grows, the baseline degrades sharply to 34.5% at 8 years, while the proposed system maintains 71.4%. The greatest improvement occurs in the 2-6 year range, most operationally relevant for long-term missing person cases.

### C. Comparison with State-of-the-Art

Table III compares the proposed system against three published approaches on the 4-year time-gap benchmark.

TABLE III: Comparison with Existing Methods (4-Year Gap)

| Method | FR | AP | RT | Acc. |
|---|---|---|---|---|
| FaceNet Only [3] | Yes | No | No | 61.2% |
| Age-cGAN+FR [5] | Yes | Yes | No | 74.8% |
| StyleGAN2+FN [6] | Yes | Yes | No | 78.3% |
| Proposed System | Yes | Yes | Yes | 84.6% |

Our system achieves the highest accuracy (84.6%) and is the only approach integrating real-time CCTV capability alongside age progression and recognition. The 6.3 percentage point gain over the best prior method is attributed to the identity-preserving loss and FAISS-based efficient retrieval.

D. Limitations

GANs cannot perfectly model individual-specific aging due to lifestyle, health, and environmental factors. Real-time processing requires dedicated GPU resources. Performance may vary across underrepresented demographic groups, and the similarity threshold may require tuning per deployment.

## VI. CONCLUSION

This paper has presented, implemented, and experimentally validated a comprehensive end-to-end system for missing person identification using predictive GAN-based age progression. Results on UTKFace and MORPH Album 2 confirm a statistically significant improvement in identification accuracy across all evaluated time gaps. Future work will explore diffusion model-based age progression, multimodal feature fusion (scars, gait), adversarial robustness, demographic fairness auditing, and optimized inference pipelines for city-wide deployment.

## REFERENCES

[1] P. Phillips et al., An Introduction to Face Recognition Technology, IEEE Intelligent Systems, vol. 15, no. 1, pp. 46-53, 2000.

[2] I. Goodfellow et al., Generative Adversarial Nets, Advances in NIPS, vol. 27, pp. 2672-2680, 2014.

[3] F. Schroff, D. Kalenichenko, and J. Philbin, FaceNet: A Unified Embedding for Face Recognition and Clustering, in Proc. IEEE CVPR, pp. 815-823, 2015.

[4] O. M. Parkhi, A. Vedaldi, and A. Zisserman, Deep Face Recognition, in Proc. BMVC, 2015.

[5] G. Antipov, M. Baccouche, and J.-L. Bernadini, Face Aging with Conditional GANs, in IEEE ICIP, pp. 2089-2093, 2017.

[6] T. Karras, S. Laine, and T. Aila, A Style-Based Generator Architecture for GANs, in Proc. IEEE/CVF CVPR, pp. 4401-4410, 2019.

[7] K. Ricanek and T. Tesafaye, MORPH: A Longitudinal Image Database of Normal Adult Age-Progression, in IEEE FG Conf., pp. 341-345, 2006.

[8] Z. Zhifei, Y. Song, and H. Qi, Age Progression/Regression by Conditional Adversarial Autoencoder, in Proc. IEEE CVPR, pp. 5810-5818, 2017.

[9] K. Zhang et al., Joint Face Detection and Alignment Using Multi-task Cascaded CNNs, IEEE Signal Process. Lett., vol. 23, no. 10, pp. 1499-1503, 2016.

[10] J. Deng et al., ArcFace: Additive Angular Margin Loss for Deep Face Recognition, in Proc. IEEE/CVF CVPR, pp. 4690-4699, 2019.

[11] Q. Cao et al., VGGFace2: A Dataset for Recognising Faces Across Pose and Age, in IEEE FG Conf., pp. 67-74, 2018.

[12] J. Johnson, A. Alahi, and L. Fei-Fei, Perceptual Losses for Real-Time Style Transfer, in ECCV, pp. 694-711, 2016.