

AHAN: Asymmetric Hierarchical Attention Network for Identical Twin Face Verification

Hoang-Nhat Nguyen

Faculty of Mathematics and Informatics, Hanoi University of Science and Technology
nhath.nh204772@sis.hust.edu.vn

Abstract

Identical twin face verification represents an extreme fine-grained recognition challenge where even state-of-the-art systems fail due to overwhelming genetic similarity. Current face recognition methods achieve over 99.8% accuracy on standard benchmarks but drop dramatically to 88.9% when distinguishing identical twins, exposing critical vulnerabilities in biometric security systems. The difficulty lies in learning features that capture subtle, non-genetic variations that uniquely identify individuals. We propose the Asymmetric Hierarchical Attention Network (AHAN), a novel architecture specifically designed for this challenge through multi-granularity facial analysis. AHAN introduces a Hierarchical Cross-Attention (HCA) module that performs multi-scale analysis on semantic facial regions, enabling specialized processing at optimal resolutions. We further propose a Facial Asymmetry Attention Module (FAAM) that learns unique biometric signatures by computing cross-attention between left and right facial halves, capturing subtle asymmetric patterns that differ even between twins. To ensure the network learns truly individuating features, we introduce Twin-Aware Pair-Wise Cross-Attention (TA-PWCA), a training-only regularization strategy that uses each subject's own twin as the hardest possible distractor. Extensive experiments on the ND_TWIN dataset demonstrate that AHAN achieves 92.3% twin verification accuracy, representing a 3.4 percentage point improvement over state-of-the-art methods.

Introduction

Consider a scenario where an unauthorized individual attempts to breach a high-security facility by exploiting their genetic similarity to an authorized twin. This represents a critical vulnerability in modern biometric systems with profound implications for national security and financial institutions. Deep learning has propelled face recognition to remarkable accuracy levels, with state-of-the-art systems achieving over 99.8% accuracy on standard benchmarks like LFW. However, these same systems experience a dramatic performance drop to approximately 88.9% when distinguishing between identical twins, revealing a fundamental limitation that could be exploited to bypass biometric authentication.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

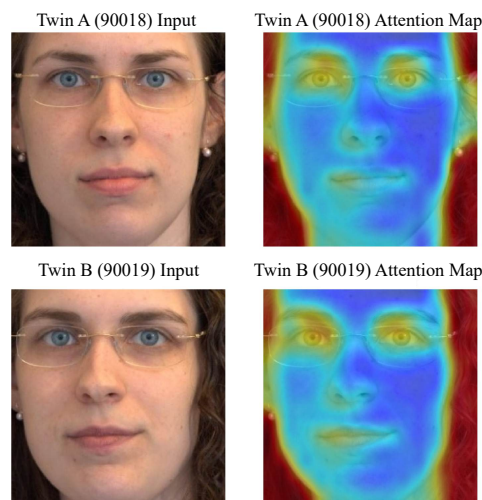


Figure 1: Visualization of the generated attention map.

The core challenge stems from twins sharing nearly 100% of their DNA, resulting in remarkably similar facial bone structure, skin texture, and overall appearance. Standard face recognition models excel at capturing global facial features like overall shape and proportions, but are easily confounded because these dominant features are nearly identical between twins. The truly identifying information lies in subtle, non-genetic characteristics: precise mole locations, unique patterns of fine lines, minor scars, slight asymmetries in facial structure, or minute differences developed over time due to environmental factors. For success in this extreme fine-grained task, models cannot rely on holistic appearance but must actively seek out and prioritize these minute, stable, highly localized differences while ignoring overwhelming genetic similarities.

Existing approaches are fundamentally ill-equipped for this challenge. State-of-the-art face recognition methods, while powerful for general populations, primarily focus on maximizing inter-class distance across diverse individuals where genetic differences create substantial feature variations. These methods lack specialized mechanisms to resolve the near-zero distance between twin identities. Con-

versely, Fine-Grained Visual Categorization (FGVC) methods like DCAL (Zhu et al. 2022) demonstrate expertise in finding subtle visual cues for distinguishing similar categories, but are general-purpose by design and lack the specific architectural priors needed for structured facial analysis.

To bridge this gap, we propose the Asymmetric Hierarchical Attention Network (AHAN), a novel framework that synergizes face recognition and fine-grained classification strengths while introducing domain-specific innovations. AHAN deconstructs twin verification into multi-granular analysis focusing precisely on features that differentiate genetically identical individuals. Our approach recognizes that different discriminative information exists at different facial locations and scales, requiring specialized analytical strategies. The key insight is that successful twin discrimination requires simultaneous analysis at three complementary levels: global facial structure (for context), local part-based features (for fine-grained details), and facial asymmetry patterns (for unique biometric signatures).

Our primary contributions are: (1) AHAN, a novel multi-stream architecture that analyzes faces at global, local, and asymmetric levels simultaneously; (2) Two new architectural modules: Hierarchical Cross-Attention (HCA) for multi-scale, part-based facial analysis and Facial Asymmetry Attention Module (FAAM) for unique asymmetry signatures; (3) Twin-Aware Pair-Wise Cross-Attention (TA-PWCA), a targeted training regularizer using each subject’s twin as the hardest distractor; (4) New state-of-the-art on ND_TWIN benchmark with 92.3% twin verification accuracy—a 3.4 percentage point improvement over existing methods.

Related Work

Highly Similar Faces Recognition

Deep face recognition has achieved significant progress through CNNs and large-scale datasets, making systems widely deployed in real-world applications. However, distinguishing between identical twins or highly similar faces remains a major challenge due to extreme similarity in facial features that often leads to higher error rates compared to general face recognition tasks.

Identical twins present the hardest cases for face recognition systems, as their facial features are nearly indistinguishable even for advanced deep learning models (Sami et al. 2022; McCauley et al. 2021; Paone et al. 2014). Studies consistently show that recognition accuracy drops significantly when distinguishing twins compared to unrelated individuals, regardless of lighting, expression, or temporal variations (Phillips et al. 2011; Paone et al. 2014; Vijayan et al. 2011). Conventional facial features and common deep learning embeddings often fail to capture the subtle differences needed to reliably separate twins (Mousavi, Charmi, and Hassanpoor 2021b,a; Paone et al. 2014).

While specialized feature extraction, hybrid models, and 3D data can improve accuracy, no method fully solves the problem in all conditions. This remains an active research area, especially for security and biometric applica-

tions where reliable twin discrimination is critical.

Fine-Grained Visual Categorization

FGVC aims to distinguish between visually similar subcategories by identifying subtle, discriminative features that differentiate closely related classes (Bi et al. 2025; Ye et al. 2024; Xie et al. 2023). Recent advances focus on automatically locating discriminative object parts without explicit annotations, with attention mechanisms becoming central to highlighting informative image regions.

Techniques such as cascaded attention and gradual attention mechanisms allow models to progressively refine focus on discriminative regions, improving classification by extracting subtle features (Xu et al. 2023; Hu et al. 2020). Methods that localize key parts using only image-level labels avoid costly part annotations while merging potential parts or using cross-layer interactions to enhance discriminative power (Zheng et al. 2020; Wang et al. 2021).

The DCAL framework (Zhu et al. 2022) represents a notable Vision Transformer-based contribution, introducing Global-Local Cross-Attention (GLCA) and Pair-Wise Cross-Attention (PWCA). GLCA addresses the limitation that standard self-attention treats all regions equally by identifying high-response regions through attention rollout and computing cross-attention between local queries and global key-value pairs. PWCA serves as training regularization using random image pairs as distractors, forcing networks to learn more robust features.

While DCAL demonstrates effectiveness for bird species and vehicle recognition, its design assumptions differ from twin verification. DCAL’s attention rollout works for clearly defined object parts but may not capture subtle facial differences. Similarly, DCAL’s random pair regularization creates moderate difficulty suitable for general fine-grained classification but insufficient for extreme twin discrimination. Our work draws architectural inspiration from DCAL’s cross-attention concepts but fundamentally re-engineers these components for facial analysis where discriminative information is more subtle and structurally constrained.

Proposed Approach

Overview

In this work, we propose the Asymmetric Hierarchical Attention Network (AHAN) to address the extreme challenge of identical twin face verification. The key insight is that distinguishing twins requires analyzing faces at multiple complementary granularities, as different types of discriminative information exist at different scales and locations. AHAN employs a multi-stream architecture built upon a Vision Transformer (Dosovitskiy 2020) backbone that simultaneously captures three types of facial representations.

The first stream uses standard self-attention to learn global facial structure and overall appearance. However, since twins share nearly identical global features due to genetic similarity, this alone is insufficient. The second stream introduces our Hierarchical Cross-Attention (HCA) module, which performs part-based analysis on semantic facial regions at multiple scales. This allows the model to capture

fine-grained details like texture patterns in the eye region or subtle shape variations in facial contours. The third stream employs our novel Facial Asymmetry Attention Module (FAAM), which learns to encode the unique asymmetric patterns that exist between the left and right halves of a face - a known biometric signature that differs even between identical twins.

To ensure the network learns truly individuating features rather than shared genetic traits, we introduce Twin-Aware Pair-Wise Cross-Attention (TA-PWCA) as a training regularization strategy. Unlike conventional methods that use random image pairs for regularization, TA-PWCA specifically uses each subject’s own twin as the hardest possible distractor during training. This forces all components of the network to focus on the subtle differences that actually distinguish twins, rather than the more obvious but genetically shared features.

Revisiting Self-Attention for Face Recognition

Before introducing our novel components, we briefly revisit the self-attention mechanism that forms the foundation of our approach. Following the standard Vision Transformer (Dosovitskiy 2020) architecture (ViT-B/16), we divide an input face image $I \in \mathbb{R}^{H \times W \times 3}$ into non-overlapping patches of size 16×16 . Each patch is linearly projected to obtain patch embeddings $\mathbf{x}_i \in \mathbb{R}^d$, where $d = 768$ is the embedding dimension and we use 12 attention heads throughout the network.

A learnable class token \mathbf{x}_{cls} is prepended to the sequence, and positional embeddings are added to encode spatial relationships:

$$\mathbf{X}_0 = [\mathbf{x}_{cls}; \mathbf{x}_1^p; \mathbf{x}_2^p; \dots; \mathbf{x}_N^p] + \mathbf{E}_{pos} \quad (1)$$

The self-attention mechanism computes attention weights by measuring the similarity between query vectors and key vectors, then uses these weights to aggregate value vectors. For a query matrix \mathbf{Q} , key matrix \mathbf{K} , and value matrix \mathbf{V} , self-attention is computed as:

$$f_{SA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right) \mathbf{V} \quad (2)$$

While self-attention excels at capturing global dependencies, it treats all facial regions equally. For twin verification, this uniform treatment is problematic because different facial regions contain different types of discriminative information that should be analyzed at different scales and with different emphasis.

Hierarchical Cross-Attention (HCA) for Part-Based Analysis

The core limitation of standard self-attention for twin verification is that it applies the same computational approach to all facial regions. However, different parts of the face contain different types of identifying information that require different analytical strategies. For example, the eye region contains rich texture information that benefits from high-resolution analysis, while the overall jawline shape is better captured at lower resolutions. Moreover, some regions

like the area around moles or scars require focused attention, while others contribute primarily to the overall facial context.

To address this challenge, we propose Hierarchical Cross-Attention (HCA), which performs structured, multi-scale analysis on semantically meaningful facial regions. Our approach is inspired by research showing that humans do not examine faces with uniform attention during identification, but rather fixate on specific regions (particularly just below the eyes) that optimize perceptual performance (Peterson and Eckstein 2012).

Semantic Region Identification: We first identify key facial landmarks using a lightweight detector to define semantic regions $\mathcal{R} = \{R_{eyes}, R_{nose}, R_{mouth}, R_{jaw}\}$. The choice of these four regions is motivated by research on human face perception: studies show that humans fixate on specific facial areas (particularly around the eyes and nose) to maximize identification performance (Peterson and Eckstein 2012), with different regions providing complementary identity information. The eye region contains rich texture (eyelashes, iris patterns); the nose provides structural geometry; the mouth captures expression-independent shape features; and the jaw defines overall facial contour. Each region R_k corresponds to a subset of patch indices that cover the relevant facial area. This semantic decomposition ensures that our attention mechanism operates on biologically-motivated facial structures rather than arbitrary image patches.

Multi-Scale Cross-Attention: For each semantic region R_k , we extract the corresponding patch embeddings and perform cross-attention at multiple scales. The key insight is that different regions benefit from different scales of analysis. We implement this through scale-adaptive queries while maintaining access to the full global context through keys and values.

For region k at scale s , we compute:

$$\mathbf{A}_k^{(s)} = \text{softmax}\left(\frac{\mathbf{Q}_k^{(s)}(\mathbf{K}^{(s)})^T}{\sqrt{d}}\right) \mathbf{V}^{(s)} \quad (3)$$

where $\mathbf{Q}_k^{(s)}$ contains region-specific queries from region k at scale s , while $\mathbf{K}^{(s)}$ and $\mathbf{V}^{(s)}$ are global representations of the entire face downsampled via average pooling at scale s . This design allows each region to attend to the full facial context while operating at its optimal resolution.

The final feature for each region is computed by aggregating across scales with learnable importance weights:

$$\mathbf{f}_k = \sum_s \alpha_{k,s} \cdot \text{AdaptivePool}(\mathbf{A}_k^{(s)}) \quad (4)$$

where $\alpha_{k,s}$ are softmax-normalized per region across scales, ensuring scale importance is learned adaptively for each facial region.

The complete HCA output concatenates features from all regions: $\mathbf{f}_{HCA} = \text{Concat}([\mathbf{f}_{eyes}, \mathbf{f}_{nose}, \mathbf{f}_{mouth}, \mathbf{f}_{jaw}])$.

This hierarchical approach allows the network to learn specialized representations for each facial region while maintaining awareness of the global facial context. The

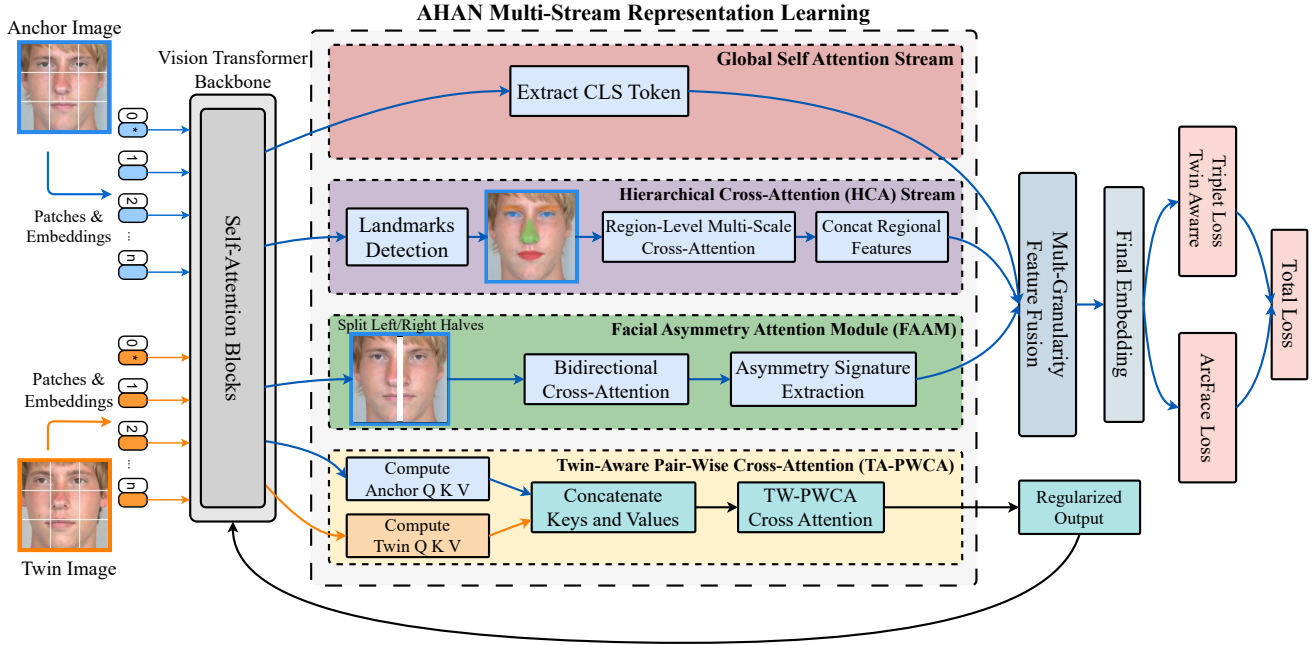


Figure 2: Overview of the AHAN architecture. The model processes face images through a Vision Transformer (Dosovitskiy 2020) backbone and three parallel streams: (1) Global Self-Attention for holistic features, (2) Hierarchical Cross-Attention (HCA) for multi-scale region analysis, and (3) Facial Asymmetry Attention Module (FAAM) for asymmetry signatures. Twin-Aware Pair-Wise Cross-Attention (TA-PWCA) uses the subject’s twin as a distractor during training only.

multi-scale design ensures that fine-grained texture details and broader structural information are both captured appropriately for each region.

Facial Asymmetry Attention Module (FAAM)

One of the most distinctive biometric characteristics that can differentiate even identical twins is facial asymmetry. While twins share nearly identical genetic facial structure, they develop unique asymmetric patterns due to environmental factors, sleeping positions, expressions, and random developmental variations. These asymmetries, though subtle, are stable over time and provide a unique biometric signature.

Traditional face recognition methods do not explicitly model facial asymmetry, instead relying on holistic representations that may inadvertently ignore these subtle but crucial differences. To address this gap, we introduce the Facial Asymmetry Attention Module (FAAM), which explicitly learns to capture and encode asymmetric patterns through cross-attention between the left and right halves of the face.

Left-Right Decomposition: We begin by spatially dividing the facial feature representation along the vertical midline. This creates two halves that should theoretically be mirror images in a perfectly symmetric face, but contain subtle differences in real faces:

$$\mathbf{X}_{left} = \mathbf{X}[:, : W/2], \mathbf{X}_{right} = \mathbf{X}[:, W/2 :] \quad (5)$$

To establish proper correspondence between the two

halves, we horizontally flip the right half so that corresponding facial landmarks align spatially:

$$\mathbf{X}_{right}^{flip} = \text{HorizontalFlip}(\mathbf{X}_{right}) \quad (6)$$

Asymmetry Cross-Attention: We compute bidirectional cross-attention between the left and right halves to capture asymmetric relationships. The attention mechanism learns to identify corresponding regions and measure their differences:

$$\mathbf{A}_{left \rightarrow right} = \text{softmax} \left(\frac{\mathbf{Q}_{left} \mathbf{K}_{right}^T}{\sqrt{d}} \right) \mathbf{V}_{right} \quad (7)$$

$$\mathbf{A}_{right \rightarrow left} = \text{softmax} \left(\frac{\mathbf{Q}_{right} \mathbf{K}_{left}^T}{\sqrt{d}} \right) \mathbf{V}_{left} \quad (8)$$

This bidirectional attention allows the model to understand how each side of the face relates to the other, capturing both correspondences and discrepancies.

Asymmetry Signature Extraction: The final asymmetry signature is computed as the absolute difference between the bidirectional attention outputs:

$$\mathbf{f}_{asym} = \text{GlobalPool}(|\mathbf{A}_{left \rightarrow right} - \mathbf{A}_{right \rightarrow left}|) \quad (9)$$

This difference vector encodes the unique asymmetric patterns of the face. Regions with high asymmetry will produce large values in this signature, while symmetric regions

will produce values close to zero. The resulting feature vector serves as a unique biometric identifier that captures the subtle asymmetric characteristics that distinguish even genetically identical individuals.

Twin-Aware Pair-Wise Cross-Attention (TA-PWCA)

While the HCA and FAAM modules provide architectural improvements for capturing discriminative facial features, standard training procedures still struggle with twin verification because the optimization naturally focuses on the most obvious differences between identities. For general face recognition, this works well because inter-class differences are substantial. However, for twins, the most obvious features are shared genetic traits, and the truly discriminative features are subtle and easily overshadowed during training.

Inspired by the Pair-Wise Cross-Attention (PWCA) concept from DCAL, we develop Twin-Aware Pair-Wise Cross-Attention (TA-PWCA) as a targeted training regularization strategy. While DCAL uses randomly sampled image pairs for regularization, we specifically use each subject’s twin as the distractor, creating the most challenging possible training scenario.

Motivation: The key insight is that if a network can successfully distinguish between twins during training, it will develop highly discriminative features that generalize well to easier verification tasks. By using twins as distractors, we force the network to ignore shared genetic features and focus exclusively on individuating characteristics.

Twin Pair Construction: During training, for each anchor image I_a of subject A , we pair it with a corresponding image I_t of subject A' (the twin of A). We compute query, key, and value matrices for both images independently.

Distraction Cross-Attention: The core of TA-PWCA involves computing cross-attention where the anchor’s queries attend to concatenated key-value pairs from both the anchor and twin images:

$$\mathbf{K}_{combined} = [\mathbf{K}_a; \mathbf{K}_t], \quad \mathbf{V}_{combined} = [\mathbf{V}_a; \mathbf{V}_t] \quad (10)$$

$$\mathbf{A}_{TA} = \text{softmax} \left(\frac{\mathbf{Q}_a \mathbf{K}_{combined}^T}{\sqrt{d}} \right) \mathbf{V}_{combined} \quad (11)$$

For each query from the anchor image, the attention mechanism must now decide whether to attend to features from the anchor itself or from its twin. This creates a challenging optimization scenario where the network must learn to distinguish between nearly identical genetic features.

Training Integration: TA-PWCA is applied during training with probability $p = 0.5$ to avoid over-regularization while ensuring sufficient exposure to twin-based distraction. The regularized attention output replaces standard self-attention in transformer layers 6-9 (the middle-to-late layers where discriminative features emerge), encouraging the network to develop twin-discriminative representations. Importantly, TA-PWCA is removed during inference, adding zero computational overhead to the deployed model.

Multi-Granularity Feature Fusion and Loss Functions

The final face embedding combines information from all three analytical streams to create a comprehensive representation that captures global structure, local part-based features, and asymmetric patterns.

Feature Fusion: We extract the global feature from the class token of the final transformer layer, then concatenate it with the HCA and FAAM outputs:

$$\mathbf{f}_{final} = \text{Concat}([\mathbf{f}_{global}; \mathbf{f}_{HCA}; \mathbf{f}_{asym}]) \quad (12)$$

This multi-granularity fusion ensures that the final embedding contains complementary information from different aspects of facial analysis.

Loss Functions: We employ a hybrid loss function that optimizes both global discrimination across all identities and fine-grained separation between twins. The ArcFace (Deng et al. 2019) loss provides strong inter-class separation across the general population, while our Twin-Aware Triplet Loss specifically addresses the twin discrimination challenge:

$$L_{total} = L_{arc} + \lambda L_{triplet} \quad (13)$$

where $\lambda = 0.1$ balances the two objectives. The twin-aware triplet loss uses cosine distance with margin $m = 0.5$ and batch-hard mining, where the negative sample is either the subject’s twin or the hardest non-twin in the batch. Twin pairs are oversampled at a 3:1 ratio to ensure sufficient exposure to the hardest discrimination cases.

Experiments

Dataset and Metrics

Dataset: We conduct experiments on the ND_TWIN dataset (Phillips et al. 2011), the most comprehensive benchmark for identical twin face verification. ND_TWIN contains 24,050 high-quality face images from 435 people at the Twins Days Festivals in Twinsburg, Ohio in 2009 and 2010, with controlled but not eliminated variations in pose, lighting, and expression. We used only frontal face images for training. The dataset is split into 87.5% training (6,336 images from 175 twin pairs), 12.5% testing (689 images from 29 twin pairs). The train/test split ensure diversity in age, race and gender in both set.

Evaluation Protocol: We follow the standard face verification protocol where positive pairs are formed by images from the same individual and negative pairs by images from different individuals. For twin-specific evaluation, we create three test scenarios with increasing difficulty: (1) *General verification*: random pairs from all test subjects following standard protocol; (2) *Twin verification*: pairs within twin families only, including same-person positive pairs (e.g., Twin A photo 1 vs Twin A photo 2) and twin-sibling negative pairs (Twin A vs Twin B); (3) *Hard twin verification*: only cross-twin pairs (Twin A vs Twin B) with no same-person comparisons—this is the hardest scenario as all pairs require distinguishing between genetically identical individuals. This multi-level evaluation provides comprehensive assessment of model performance across different difficulty levels.

Component	Setting	Value
Optimizer	Adam	lr=1e-4, wd=5e-4
Schedule	Cosine	100 epochs
Batch Size	64	4 accum steps
HCA Regions	4	eyes, nose, mouth, jaw
HCA Scales	3	1×, 2×, 4×
TA-PWCA Prob	0.5	Training only
Loss Weight λ	0.1	ArcFace + Triplet

Table 1: Key hyperparameters and training configuration.

Metrics: We report multiple metrics to provide thorough evaluation: Verification Accuracy (Acc), False Acceptance Rate at 0.1% and 0.01% (FAR@0.1%, FAR@0.01%), True Accept Rate at 1% False Accept Rate (TAR@1%FAR), and ROC Area Under Curve (AUC). These metrics capture both overall performance and behavior at critical operating points for security applications.

Implementation Details

Training Setup: We implement AHAN using PyTorch on 1 NVIDIA P100 GPU with 16GB memory each. The ViT-B/16 backbone is initialized with ImageNet-21k pre-trained weights. All face images are preprocessed by MTCNN for detection and alignment, then resized to 224×224 pixels. We apply standard data augmentation including random horizontal flipping, color jittering (brightness ± 0.2 , contrast ± 0.2 , saturation ± 0.2), and random rotation (± 10).

Hyperparameters: Training uses Adam optimizer with initial learning rate $1e-4$, weight decay $5e-4$, and cosine annealing schedule over 100 epochs. Batch size is set to 64 with gradient accumulation over 4 steps. For HCA, we use 4 semantic regions (eyes, nose, mouth, jaw) with 3 scales each (1×, 2×, 4× downsampling). TA-PWCA is applied with probability 0.5 during training. The combined loss uses $\lambda = 0.1$ for triplet loss weighting. The facial landmark detector is MediaPipe with lightweight MobileNet backbone.

Baselines

We compare AHAN against strong baselines from **State-of-the-Art Face Recognition**: (1) ArcFace (Deng et al. 2019) with ResNet-100 backbone, (2) CosFace (Wang et al. 2018) with ResNet-100, (3) AdaFace (Kim, Jain, and Liu 2022) with IR-101, (4) MagFace (Meng et al. 2021) with IR-100, (5) TransFace (Sun and Tzimiropoulos 2022) with ViT-B/16, (6) UniFace (Zhou et al. 2023) with unified cross-entropy loss, and (7) CurricularFace (Huang et al. 2020) with curriculum learning. We also include **Fine-Grained Methods**: (8) TransFG (He et al. 2022) for discriminative token selection. These represent current best practices in face recognition and fine-grained classification.

Main Results

Table 3 presents comprehensive results comparing AHAN with all baseline methods across three evaluation scenarios. AHAN achieves significant improvements over all baselines, with particularly strong performance in the challenging twin verification scenarios.

Method	Backbone	Key Innovation
ArcFace	ResNet-100	Angular margin
CosFace	ResNet-100	Cosine margin
AdaFace	IR-101	Adaptive margin
CurricularFace	IR-101	Curriculum learning
UniFace	ResNet-100	Unified cross-entropy
TransFace	ViT-B/16	Vision transformer
TransFG	ViT-B/16	Discriminative tokens
AHAN (Our)	ViT-B/16	Multi-granularity + asymmetry

Table 2: Comparison with baseline methods

The results demonstrate that AHAN significantly outperforms all baselines across all metrics and scenarios. For twin verification, AHAN achieves 92.3% accuracy compared to the best baseline (ArcFace) at 88.9%, representing a 3.4 percentage point improvement. In the experimental hard twin verification scenario, AHAN achieves 88.5% accuracy versus ArcFace’s 85.3%, a 3.2 percentage point improvement. This validates the effectiveness of our multi-granularity approach for extreme fine-grained face verification.

Qualitative Analysis

Attention Visualization: Figure 1 presents attention maps between an identical twin pair. AHAN successfully focuses on discriminative facial regions like subtle texture around the eyes and unique patterns near the mouth.

Failure Case Analysis: We analyze failure cases where AHAN incorrectly verifies twin pairs. Common failure modes include: (1) Extreme lighting variations, (2) Time capture changing (images capture in 2009 and 2010). These cases represent fundamental limits of appearance-based verification and suggest future directions for improvement.

Ablation Studies

We conduct extensive ablation studies to validate the contribution of each component. Table 4 shows that each module contributes significantly to the final performance, with HCA providing the largest individual improvement (+5.9%) and TA-PWCA offering substantial training benefits (+8.7%).

Component Analysis: The results demonstrate that each component contributes meaningfully to performance. HCA provides the largest single improvement (+5.9% on hard twin verification), followed by FAAM (+4.5%). The combination of HCA and FAAM is synergistic, achieving +10.1% improvement together. TA-PWCA regularization adds an additional +4.3% improvement, validating the importance of twin-aware training.

Design Choices for HCA: We analyze different design choices for the Hierarchical Cross-Attention module. The optimal configuration uses 4 regions (eyes, nose, mouth, jaw), balancing performance and computational cost. Three scales provide optimal performance, capturing fine details (1×), intermediate features (2×), and global context (4×).

Method	General Verification			Twin Verification			Hard Twin Verification		
	Acc	AUC	TAR@1%FAR	Acc	AUC	TAR@1%FAR	Acc	AUC	TAR@1%FAR
Non Transformer-based Face Recognition									
ArcFace	98.5	99.4	95.6	88.9	93.8	82.4	85.3	90.6	78.4
MagFace	98.4	99.3	95.4	88.5	93.4	81.9	85.0	90.3	78.0
AdaFace	98.2	99.2	95.1	88.2	93.1	81.5	84.7	90.0	77.6
CosFace	98.0	99.1	94.8	87.5	92.5	80.6	84.1	89.4	76.8
CurricularFace	98.3	99.3	95.2	87.1	92.1	80.1	83.6	88.9	76.2
UniFace	97.9	99.0	94.6	86.7	91.8	79.6	83.2	88.5	75.7
Transformer-based Methods									
TransFace	97.5	98.7	94.0	85.2	90.4	77.8	81.8	87.2	73.9
TransFG	97.3	98.5	93.6	84.8	90.0	77.3	81.4	86.8	73.4
Ours									
AHAN	99.1	99.8	97.2	92.3	96.4	87.6	88.5	93.5	82.8

Table 3: Main results on ND_TWIN dataset. Best results in bold.

Configuration	TAR@		
	FAR=0.1%	EER (%)	TVA (%)
Baseline (ViT-B)	52.1	24.7	81.2
+ HCA	63.4	19.3	87.1
+ FAAM	58.9	21.8	84.7
+ TA-PWCA	67.8	17.1	89.9
+ HCA + FAAM	69.2	16.4	90.4
+ HCA + TA-PWCA	74.6	14.2	91.8
+ FAAM + TA-PWCA	71.3	15.7	90.7
Full AHAN	78.4	12.1	92.3

Table 4: Ablation study results on Hard Twin Verification.

MediaPipe provides the best balance of accuracy, speed, and final performance for landmark detection.

Regularization Analysis: We compare TA-PWCA with alternative regularization strategies. TA-PWCA significantly outperforms all alternatives including random pairs (original PWCA), hard negative mining, and sibling pairs, confirming that using twins as distractors provides the most effective regularization for this task.

Efficiency Analysis: AHAN requires 33% more parameters and 36% more FLOPs than the baseline ViT, but provides substantial performance gains (+6.7 percentage points on hard twin verification compared to TransFace). The computational overhead is reasonable for high-security applications where accuracy is paramount.

Discussion

Component Synergy: Our ablation studies demonstrate that each AHAN component addresses specific limitations in twin verification. HCA enables multi-scale regional analysis where eye texture benefits from fine-grained processing while jawline shape requires broader context. FAAM captures asymmetry signatures—a biometric characteristic that distinguishes even identical twins due to environmental factors and developmental variations. TA-PWCA regularization

forces the network to ignore shared genetic features and focus on truly individuating characteristics. The synergistic combination creates a system specifically engineered for extreme fine-grained verification, achieving improvements that no single component could provide alone.

Limitations and Future Work: AHAN shows robustness but has limitations in extreme scenarios. Performance degrades with severe pose variations ($> 45^\circ$), extensive occlusion ($> 40\%$), and large temporal gaps (> 5 years). The method requires facial landmark detection which can fail in extreme conditions; future work could explore landmark-free approaches. Incorporating other modalities (gait, voice) could further improve discrimination, and developing more efficient architectures remains important for resource-constrained deployment. Privacy implications of reliable twin verification must also be carefully considered.

Conclusion

This paper presents the Asymmetric Hierarchical Attention Network (AHAN), a novel architecture for identical twin face verification. Our key innovation is that twin discrimination requires simultaneous analysis at multiple granularities: global facial structure, local part-based features, and facial asymmetry patterns. Through Hierarchical Cross-Attention (HCA) for multi-scale part analysis, the Facial Asymmetry Attention Module (FAAM), and Twin-Aware Pair-Wise Cross-Attention (TA-PWCA) regularization, AHAN achieves 92.3% twin verification accuracy—a 3.4 percentage point improvement over ArcFace.

The success of AHAN establishes a new paradigm for extreme fine-grained biometric verification, demonstrating that domain-specific architectural innovations can overcome challenges that generic approaches cannot address. This work has significant implications for high-security biometric applications and opens promising research directions including extension to other fine-grained recognition tasks and integration with complementary biometric modalities.

References

- Bi, Q.; Zhou, B.; Ji, W.; and Xia, G.-S. 2025. Universal fine-grained visual categorization by concept guided learning. *IEEE Transactions on Image Processing*.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4690–4699.
- Dosovitskiy, A. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- He, J.; Chen, J.-N.; Liu, S.; Kortylewski, A.; Yang, C.; Bai, Y.; and Wang, C. 2022. Transfg: A transformer architecture for fine-grained recognition. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 852–860.
- Hu, Y.; Yang, Y.; Zhang, J.; Cao, X.; and Zhen, X. 2020. Attentional kernel encoding networks for fine-grained visual categorization. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(1): 301–314.
- Huang, Y.; Wang, Y.; Tai, Y.; Liu, X.; Shen, P.; Li, S.; Li, J.; and Huang, F. 2020. Curricularface: adaptive curriculum learning loss for deep face recognition. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5901–5910.
- Kim, M.; Jain, A. K.; and Liu, X. 2022. Adaface: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18750–18759.
- McCauley, J.; Soleymani, S.; Williams, B.; Dando, J.; Nasrabadi, N.; and Dawson, J. 2021. Identical twins as a facial similarity benchmark for human facial recognition. In *2021 International Conference of the Biometrics Special Interest Group (BIOSIG)*, 1–5. IEEE.
- Meng, Q.; Zhao, S.; Huang, Z.; and Zhou, F. 2021. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14225–14234.
- Mousavi, S.; Charmi, M.; and Hassanpoor, H. 2021a. A distinctive landmark-based face recognition system for identical twins by extracting novel weighted features. *Computers & Electrical Engineering*, 94: 107326.
- Mousavi, S.; Charmi, M.; and Hassanpoor, H. 2021b. Recognition of identical twins based on the most distinctive region of the face: Human criteria and machine processing approaches. *Multimedia Tools and Applications*, 80(10): 15765–15802.
- Paone, J. R.; Flynn, P. J.; Philips, P. J.; Bowyer, K. W.; Bruegge, R. W. V.; Grother, P. J.; Quinn, G. W.; Pruitt, M. T.; and Grant, J. M. 2014. Double trouble: Differentiating identical twins by face recognition. *IEEE Transactions on Information Forensics and Security*, 9(2): 285–295.
- Peterson, M. F.; and Eckstein, M. P. 2012. Looking just below the eyes is optimal across face recognition tasks. *Proceedings of the National Academy of Sciences*, 109(48): E3314–E3323.
- Phillips, P. J.; Flynn, P. J.; Bowyer, K. W.; Bruegge, R. W. V.; Grother, P. J.; Quinn, G. W.; and Pruitt, M. 2011. Distinguishing identical twins by face recognition. In *2011 IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, 185–192. IEEE.
- Sami, S. M.; McCauley, J.; Soleymani, S.; Nasrabadi, N.; and Dawson, J. 2022. Benchmarking human face similarity using identical twins. *IET Biometrics*, 11(5): 459–484.
- Sun, Z.; and Tzimiropoulos, G. 2022. Part-based face recognition with vision transformers. *arXiv preprint arXiv:2212.00057*.
- Vijayan, V.; Bowyer, K. W.; Flynn, P. J.; Huang, D.; Chen, L.; Hansen, M.; Ocegueda, O.; Shah, S. K.; and Kakadiaris, I. A. 2011. Twins 3D face recognition challenge. In *2011 international joint conference on biometrics (IJCB)*, 1–7. IEEE.
- Wang, H.; Wang, Y.; Zhou, Z.; Ji, X.; Gong, D.; Zhou, J.; Li, Z.; and Liu, W. 2018. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5265–5274.
- Wang, J.; Li, N.; Luo, Z.; Zhong, Z.; and Li, S. 2021. High-order-interaction for weakly supervised fine-grained visual categorization. *Neurocomputing*, 464: 27–36.
- Xie, Y.; Gong, Q.; Luan, X.; Yan, J.; and Zhang, J. 2023. A survey of fine-grained visual categorization based on deep learning. *Journal of Systems Engineering and Electronics*, 35(6): 1337–1356.
- Xu, Q.; Zhang, M.; Li, Y.; and Tao, Z. 2023. Learning more discriminative clues with gradual attention for fine-grained visual categorization. *Image and Vision Computing*, 136: 104753.
- Ye, S.; Yu, S.; Wang, Y.; and You, X. 2024. R2-trans: Fine-grained visual categorization with redundancy reduction. *Image and Vision Computing*, 143: 104923.
- Zheng, X.; Qi, L.; Ren, Y.; and Lu, X. 2020. Fine-grained visual categorization by localizing object parts with single image. *IEEE Transactions on Multimedia*, 23: 1187–1199.
- Zhou, J.; Jia, X.; Li, Q.; Shen, L.; and Duan, J. 2023. Uni-face: Unified cross-entropy loss for deep face recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 20730–20739.
- Zhu, H.; Ke, W.; Li, D.; Liu, J.; Tian, L.; and Shan, Y. 2022. Dual cross-attention learning for fine-grained visual categorization and object re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4692–4702.