

# FusionAgent: A Multimodal Agent with Dynamic Model Selection for Human Recognition

Jie Zhu<sup>1</sup> Xiao Guo<sup>1</sup> Yiyang Su<sup>1</sup> Anil Jain<sup>1</sup> Xiaoming Liu<sup>1,2</sup>  
<sup>1</sup>Michigan State University <sup>2</sup>University of North Carolina at Chapel Hill  
{zhujie4, guoxia11, suyiyang1, jain}@msu.edu liuxm@cs.unc.edu

## Abstract

Model fusion is a key strategy for robust recognition in unconstrained scenarios, as different models provide complementary strengths. This is especially important for whole-body human recognition, where biometric cues such as face, gait, and body shape vary across samples and are typically integrated via score-fusion. However, existing score-fusion strategies are usually static, invoking all models for every test sample regardless of sample quality or modality reliability. To overcome these limitations, we propose **FusionAgent**, a novel agentic framework that leverages a Multimodal Large Language Model (MLLM) to perform dynamic, sample-specific model selection. Each expert model is treated as a tool, and through Reinforcement Fine-Tuning (RFT) with a metric-based reward, the agent learns to adaptively determine the optimal model combination for each test input. To address the model score misalignment and embedding heterogeneity, we introduce Anchor-based Confidence Top-k (ACT) score-fusion, which anchors on the most confident model and integrates complementary predictions in a confidence-aware manner. Extensive experiments on multiple whole-body biometric benchmarks demonstrate that FusionAgent significantly outperforms SoTA methods while achieving higher efficiency through fewer model invocations, underscoring the critical role of dynamic, explainable, and robust model fusion in real-world recognition systems. Project page: [FusionAgent](#).

## 1. Introduction

Computer vision has made remarkable progress through state-of-the-art models for tasks such as object detection, segmentation, face recognition, and image classification [4, 12, 14, 18, 25, 26, 65, 66]. Yet existing research has largely focused on single-model performance, paying far less attention to how complementary models can be effectively combined. Since different models often excel under different conditions, selective fusion offers a promising path toward

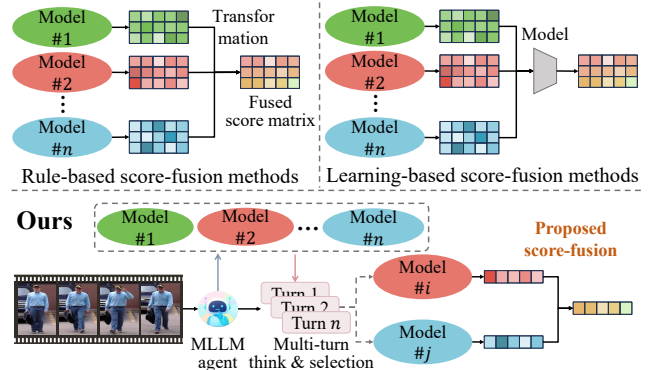


Figure 1. **Comparison of score-fusion methods.** **Top:** Rule-based methods apply predefined transformations to fuse all model scores, while learning-based methods infer a fusion model from data but still assume that every model contributes to all test samples. **Bottom:** our framework leverages an MLLM agent to dynamically select a subset of models, followed by the proposed score-fusion strategy, enabling adaptive and robust integration.

stronger and more robust performance.

Model fusion is widely used with applications ranging from feature fusion to score fusion [46, 63, 64]. Among these, score fusion is particularly representative and practical, especially in biometric recognition [19, 35, 67]. As shown in Fig. 1, existing methods are mainly rule-based or learning-based. The former combine scores with predefined rules [19, 46], while the latter learn fusion strategies from data [35, 50, 67]. Despite their differences, both typically rely on a fixed combination of models for all inputs, assuming universal complementarity. This assumption is often violated in practice: for example, face recognition scores are unreliable when only a person’s back view is visible. Even methods such as QME [67], which account for input quality, still allow low-content inputs to affect the final output. These limitations motivate a more selective fusion paradigm that first determines which models are suitable for each input, and then how to combine their outputs effectively. To systematically study this problem, we pose the central research question of this paper:

*How can we adaptively choose the **optimal model combination for each test sample** while rigorously testing whether such selection improves performance even with an **ad hoc score-fusion algorithm**?*

This raises the question of how to perform sample-wise model selection in a principled and interpretable manner. Recent advances in multimodal large language models (MLLMs) suggest a promising direction for adaptive model fusion [3, 13, 33, 55, 69]. By reasoning over complex inputs and orchestrating external tools in a context-dependent manner, an agent can treat each specialized model as a tool and decide which ones to invoke for each sample. This makes agentic fusion particularly suitable for selective fusion, where model reliability varies across inputs and decisions must be made adaptively for each sample.

This motivation is especially compelling in systematic whole-body recognition [67], which integrates multimodal traits (e.g., face, gait, and body shape) and multiple specialized models. Unlike unimodal systems based on face recognition (FR) [8, 22, 24], gait recognition (GR) [56, 61], or person re-identification (ReID) [11, 29, 48], whole-body recognition depends on complementary modalities whose reliability varies with input conditions. For example, face cues may fail under occlusion or extreme viewpoints, while gait or body appearance remains informative [28]. This makes whole-body recognition a representative testbed for agentic model selection and fusion.

Inspired by these, we introduce a unified framework centered around two core components, each designed to systematically respond to model selections and score-fusion methods. In response to model selections, we propose **FusionAgent**, an MLLM-based agent that performs dynamic model selection per sample. Each biometric model is wrapped as a tool, providing the agent with a score vector and predicted label. Through Reinforcement Fine-Tuning (RFT) guided by the **proposed metric-based reward**, the agent learns to identify sample-specific model combinations by analyzing query patterns, effectively answering how optimal ensembles can be found for each input. To explore score-fusion methods, we design a lightweight **Anchor-based Confidence Top-k score-fusion method (ACT)**. This method leverages the most confident model selected by the agent, then integrates only the top-matching scores scaled by confidence weights. By doing so, ACT mitigates misalignment from sample-wise model selection and embedding heterogeneity, demonstrating that even a simple fusion rule can yield strong performance when guided by adaptive model selection. Our main contributions are:

- We propose FusionAgent, an agentic framework that leverages an MLLM to perform explainable, sample-wise dynamic model selection.
- We introduce an Anchor-based Confidence Top-k score-

fusion method (ACT) to mitigate score misalignment arising from dynamic model selection, ensuring robust integration of heterogeneous score outputs.

- We design a metric-based reward function, central to our optimization, which directly aligns the agent’s selection strategy with final performance metrics.
- Extensive experiments on multiple whole-body biometric benchmarks demonstrate the superiority of our approach over state-of-the-art methods, even using conventional score-fusion methods.

## 2. Related Work

### 2.1. Whole-body Biometric Recognition

Whole-body biometric recognition systems combine detectors (e.g., for whole body and face), embedding models, and fusion modules to leverage multi-modal cues such as face and body features [7, 28, 30, 49]. The key challenge is to effectively integrate the complementary strengths of different modalities and their dedicated models to maximize overall performance. For example, FR models excel with high-quality frontal faces but struggle under adverse conditions like oblique angles [9, 22–24]. In contrast, GR models focus on clothing-invariant dynamic body attributes [18, 56, 57, 61, 62], while ReID models take a holistic approach to extract comprehensive appearance features [11, 27, 29, 52, 54]. Prior fusion methods typically use all available models for every query [5, 21, 28, 30, 35, 50, 67], ignoring the sample-dependent nature of optimal model combinations. A clear case is a low-resolution image with a side-view of a person: FR models may be unreliable, making GR or ReID models more critical for that specific sample. Therefore, we propose FusionAgent to dynamically select an optimal model combination for each test sample through explainable analysis, thereby tailoring the fusion process to individual inputs.

### 2.2. Score-fusion

Unlike feature-level fusion [1, 25, 36, 37, 48, 63], score-level fusion integrates similarity scores from multiple modalities to improve recognition performance [45]. This approach is broadly categorized into two paradigms: rule-based and learning-based methods. Rule-based methods employ fixed rules such as Z-score, Min-max normalization, max/min fusion [19, 20, 40, 58], and likelihood ratio-based fusion [15, 34, 37, 38, 51], offering simplicity and efficiency. In contrast, learning-based methods optimize fusion during training [5, 35, 50]. These include the recent quality-aware approach QME [67], which performs weighted score-fusion based on input quality. In contrast, we introduce a simple score-fusion algorithm with an anchor. This design, combined with selective model inputs, significantly outperforms existing fusion strategies.

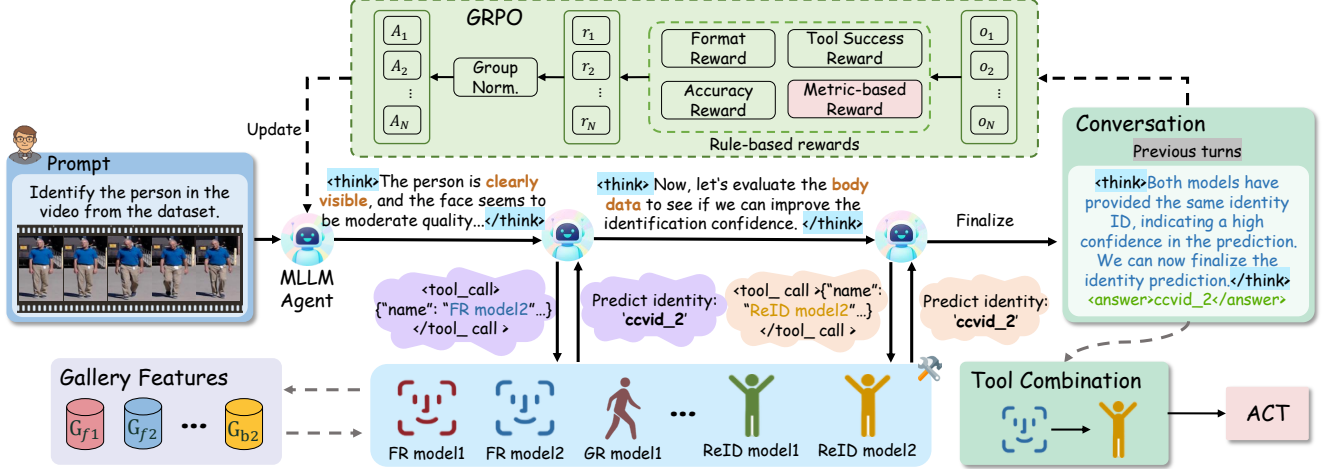


Figure 2. **Overview of the FusionAgent framework.** Recognition models are wrapped as tools to generate score vectors and predicted identities based on gallery features. The MLLM agent receives multimodal biometric inputs and performs a reasoning-action step through multi-turn, selectively invokes tools, and integrates predictions into a final identity decision and fused score vector. The agent is optimized with reinforcement fine-tuning using rule-based rewards, including the proposed metric-based reward.

### 2.3. Agents and Reinforcement Learning

The advent of tool-augmented MLLMs has spurred the development of agentic systems: models capable of planning, reasoning, and interacting with external tools to solve complex tasks. Recent studies [3, 13, 33, 55, 69] have equipped MLLMs with tools to address diverse challenges such as visual grounding, image super-resolution, and advanced visual understanding. Furthermore, subsequent research has employed RFT with rule-based rewards, such as Group Relative Policy Optimization (GRPO) [41], which offers improved generalizability with lower data dependency [32, 42, 53, 59, 68]. Building upon these advancements, we further propose a metric-based reward to optimize model selections of FusionAgent, effectively capturing data-dependent patterns.

## 3. Proposed Method

**Problem Formulation.** In biometric recognition, a query (or probe) is a sample sequence that is to be identified (1:N comparisons) or verified by comparison (1:1 comparison) with a gallery of previously enrolled subjects in the system. Let  $Q$  denote the set of query videos and  $G$  be the gallery.  $M$  is a predefined model set  $\{m_1, m_2, \dots, m_Z\}$  that is used for feature extraction. For a query  $q \in Q$  and a model  $m \in M$  (from a predefined model set  $M$ ),  $m$  extracts features from  $q$  and produces a similarity score vector  $s_{m,q} \in \mathbb{R}^{|G|}$  with respect to each gallery entry  $g \in G$ . The human recognition pipeline comprises two stages: (i) *sample-level model subset selection* and (ii) *query-based score fusion of selected models*. For each  $q$ , we seek an optimal subset of models  $M_q \subseteq M$  that maximizes per-

formance. Given  $M_q$ , we stack the per-model scores into a score matrix  $S_q \in \mathbb{R}^{|M_q| \times |G|}$  whose  $i$ -th row is  $s_{m_i,q}$ . A score-fusion function  $f$  then maps  $S_q$  to a fused score vector  $s'_q = f(S_q) \in \mathbb{R}^{|G|}$ . We detail the model-selection procedure using an agent in Sec. 3.1, the supervision signals (reward functions) in Sec. 3.2, and the proposed score-fusion method in Sec. 3.3.

### 3.1. Agentic Training Framework

Training the agent presents two challenges: (i) the lack of ground-truth dialogue data for Supervised Fine-tuning (SFT), since grid searching optimal model combinations per sample is infeasible; and (ii) the intractability of dataset-level grid search for the best model combination per sample due to exponential growth in the search space with models and samples. We thus employ GRPO, which enables trial-based learning without pre-labeled data. Additionally, we design a metric-based reward function that promotes exploration of diverse model combinations and facilitates implicit learning of the relationship between input characteristics and model performance, as detailed in Sec. 3.2.

The overall framework is depicted in Fig. 2. For a given query  $q$ , the agent first analyzes its content and selects an initial model  $a_q \in M_q$ . For example, FR model 2 is selected because a clearly visible face is detected. Upon a successful function call, the tool returns  $s_{a_q}(q)$  and the predicted label. Since score vectors are not directly interpretable by the agent, the predicted identity label is returned explicitly, while  $s_{a_q}(q)$  is retained internally. The agent then decides whether to invoke another model or output its decision. This sequential decision-making process enables the flexible maximization of a long-term objective. [44, 55]. If

the agent decides to output its decision, it also provides a reasoning summary that justifies the model selections. Reasoning at each step is necessary to enhance transparency and decision traceability. The whole conversation  $o_i$ , including the successful model selections, is then used to measure the overall rewards. In GRPO, we sample  $N$  rollouts (*i.e.*, responses) for each query and jointly compute their advantages  $\{A_1, A_2, \dots, A_N\}$  from the corresponding rewards  $\{r_1, r_2, \dots, r_N\}$  to update the agent. Details of GRPO are provided in the supplementary.

**Multi-turn Design.** We adopt a ReAct-style (reason-before-act) multi-turn controller for tool use, rather than a single-shot plan generator [55]. This design is motivated by the need to decompose the complex task of multi-tool usage into a sequence of simpler, more manageable decisions. A single-turn approach would require the agent to generate a complete and static execution plan at once, a task plagued by a combinatorial action space and an inability to handle unexpected outcomes. This interleaving of reasoning with actions decomposes multi-tool execution into atomic steps and yields the following benefits:

- i) **Simplified Learning:** Reduces the vast action space of generating a full plan to a single decision at each step, making the policy significantly more tractable to learn.
- ii) **Dynamic Adaptation:** Allows the agent to observe tool outputs and adjust its strategy in real time, enabling error correction and flexible reasoning. These sequential actions also support effective credit assignment during inference.

### 3.2. Reward Functions

Reward functions play a central role in RFT-based agent training. We design four reward functions: format reward  $R_f$ , tool success reward  $R_{tool}$ , answer accuracy reward  $R_{acc}$ , and metric-based reward  $R_{mat}$ .

**Format Reward.** The agent is trained to produce structured responses that separate reasoning, tool calls, and final answers [32, 33, 41, 68]. We adapt this to a multi-turn setting where each turn must be a structured format, such as `<think>...</think> <answer>...</answer>`. The reward is computed per turn, and the overall reward is averaged across all turns.

**Tool Success Reward.** This reward assesses whether the agent’s tool calls are executable and yield valid results. Each tool call receives a binary success/failure score, and the reward is defined as the success rate across all tool calls in a trajectory. This incentivizes the agent to produce syntactically correct tool inputs.

**Answer Accuracy Reward.** This reward measures the correctness of the agent’s final prediction  $a_i$  extracted from the answer tag `<answer>...</answer>` of the response  $o_i$  against the ground truth label  $y$ . Its primary objective is to prioritize factual accuracy over procedural correctness,

which enables the agent to implicitly assess model reliability, especially when model predictions are inconsistent:

$$R_{acc}(a_i, y) = \begin{cases} 1, & \text{if } a_i = y, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

**Metric-based Reward.** The metric-based reward guides the agent toward effective and dynamic model selection. A key challenge is to encourage sufficient exploration of diverse model combinations without prior knowledge of their optimality. To address this, the reward is computed based on the performance of the agent-selected model combination on the training set, using the score-fusion method from Sec. 3.3. The combination is derived from successfully executed conversations per sample.

However, exhaustively searching for the ground-truth optimal model combination for each sample is computationally infeasible. To promote exploration, we construct an augmented model selection  $\mathbf{M}_Q \in \mathbb{R}^{|Q| \times |M|}$  based on the model combination  $M_{o_i}$  of the response  $o_i$ , since TAR and FNIR are determined by thresholds selected from *dataset-level* scores. Specifically, a fraction  $\gamma \in [0, 1]$  of samples in  $\mathbf{M}_Q$  retain the same model combination as  $M_{o_i}$ . For the remaining  $(1 - \gamma)$ , the mask entries are sampled from a Bernoulli distribution to diversify the model combinations explored. The column corresponding to the anchor model is always set to true, ensuring its consistent inclusion.

We apply the ACT (Anchor-based Confidence Top-k) (details in Sec. 3.3) based on  $\mathbf{M}_Q$ . The performance of the resulting fused scores matrix  $\mathbf{S}'_{o_i}$  is then evaluated across the entire training set using four key metrics: True Accept Rate at a False Acceptance Rate (TAR@FAR), mean Average Precision (mAP), Rank-1 accuracy, and False Non-Identity Rate at a specified False Positive Identification Rate (FNIR@FPIR). The final metric-based reward  $R_{mat}$  is a composite score formulated to provide a holistic performance signal:

$$\mathbf{S}'_{o_i} = \text{ACT}(\mathbf{M}_Q, k), \quad (2)$$

$$R_{mat}(o_i) = \text{Rank}(\mathbf{S}'_{o_i}) + \text{mAP}(\mathbf{S}'_{o_i}) + \text{TAR}(\mathbf{S}'_{o_i}) - \text{FNIR}(\mathbf{S}'_{o_i}). \quad (3)$$

This formulation naturally aligns with both our optimization objective and real-world deployment needs. It rewards the agent for achieving higher accuracy and retrieval performance (TAR, mAP, Rank) while penalizing missed identifications (FNIR), thereby capturing operational performance requirements in a comprehensive manner. The final reward is the sum of  $R_f$ ,  $R_{tool}$ ,  $R_{acc}$ , and  $R_{mat}$ .

### 3.3. Anchor-based Confidence Top-k Score-fusion

The proposed ACT score-fusion approach aims to dynamically and effectively combine scores from multiple models for a given query. The overview is shown in Fig. 3.




Figure 3. **Overview of the ACT score-fusion.** Based on tool execution results (*i.e.*, score vectors) and selected model combination, the first selected model serves as the anchor, and ACT produces the final score vector via confidence weighting and top-k filtering.

Our approach is built on the principle of leveraging a stable and powerful “anchor model”  $m_a \in \mathbf{M}_q$  to provide a robust score vector, while selectively incorporating normalized, high-confidence scores from the set of models. Let  $m_a$  be the first selected model in  $\mathbf{M}_q$ , and  $s_{m,q,g}$  denotes the similarity score between  $q$  and  $g$  for model  $m$ .

We begin by computing a contribution score  $c_{m,q,g}$ . This score is designed to leverage the most confident predictions while filtering out potential noise from low-scoring candidates, which is achieved through a Top-k selection process. Formally, for a score  $c_{m,q,g}$  in  $\mathbf{c}_{m,q} \in \mathbb{R}^{|\mathcal{G}|}$ , the contribution score is defined as:

$$c_{m,q,g} = \begin{cases} z_{m,q,g} \cdot s_{m,q,g} & \text{if } g \in \mathcal{T}_{m,q}, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

where  $\mathcal{T}_{m,q}$  is the set of indices of the  $k$  highest-scoring gallery entries for query  $q$  from model  $m$ . The term  $z_{m,q,g}$  represents the Z-score normalized score between  $q$  and  $g$  by model  $m$ . This standardization makes the scores from different models with potentially different scales comparable. It is calculated as:

$$z_{m,q,g} = \frac{s_{m,q,g} - \mu_{m,q}}{\sigma_{m,q}}, \quad (5)$$

where  $\mu_{m,q}$  and  $\sigma_{m,q}$  are the mean and standard deviation of  $s_{m,q}$ . The final fused score vector  $\mathbf{s}'_q$  is computed via:

$$\mathbf{s}'_q = \frac{1}{1 + |\mathbf{M}_q|} \left( \mathbf{s}_{m_a,q} + \sum_{m \in \mathbf{M}_q} \mathbf{c}_{m,q} \right). \quad (6)$$

The term  $(1 + |\mathbf{M}_q|)$  serves as the normalization factor, ensuring a balanced contribution from all models. This strategy balances robust, general performance with specialized, high-confidence insights. The anchor model provides a more substantial contribution by establishing a global ranking structure, as its scores are applied unconditionally to all candidates. In contrast, the selected models provide sparse, localized refinements only for their top-k predictions. Fig. 4 provides a toy example of ACT score-fusion.




Figure 4. **A toy example of the proposed ACT score-fusion.** Three models are used with the FR model as the anchor and  $k = 1$ . ACT amplifies the gap between match and non-match scores through confidence-based top-k and anchor weighting, improving verification and open-set search performance.

## 4. Experiments

**Datasets and Evaluation Protocol.** We evaluate on popular human recognition datasets: CCVID [11], MEVID [6], and LTCC [39]. They comprise multi-view captures and cross-modal biometric data with diverse resolutions and temporal dynamics. Our evaluation protocol follows [67] for fair comparison. Performance is measured via standard metrics in general settings, *i.e.*, Rank-1 accuracy, mean Average Precision (mAP), verification rate (TAR@1.0%FAR), and open-set search metric (FNIR@1.0%FPIR), which collectively reflect real-world deployment requirements. For open-set search evaluation, we follow the protocol of [47, 67] to construct 10 random subsets of gallery subjects. Each subset, containing  $\sim 20\%$  of the subjects in the test set, serves as a non-mated list. We report the mean and standard deviation for the 10 trials.

**Baselines.** Our baselines include the following methods: Min/Max score-fusion [19], Z-score and min-max normalization [46], RHE [15], FarSight [28], and learning-based methods like BSSF [50], Weighted-sum [35], AsymA-O1 [16], and QME [67]. We follow [67] to construct the biometric model pool (details provided in the Supp.). We also compare with the SoTA multimodal human recognition model SapiensID [25].

**Implementation Details.** Following [67], we extract the center embeddings of subjects for training and gallery embeddings for testing by applying average pooling. Considering runtime efficiency in practical scenarios, Fusion-Agent is based on Qwen2.5-VL-3B [2] and fine-tuned via GRPO [41]. For efficiency, we set a turn limit of 4 and train for 200 steps. All biometric models are frozen in training. For GRPO, we set  $N = 6$  and a KL coefficient  $\beta = 0.04$ . LoRA [17] is applied with rank  $r = 64$  and scaling factor  $\alpha = 128$ . The learning rate follows a linear decay schedule starting from  $2 \times 10^{-5}$ . The metric-based reward uses  $\gamma = 0.8$ . Based on training-set performance, we set  $k = 10$

<i>Method</i>	<i>Comb.</i>	Rank1 $\uparrow$	mAP $\uparrow$	TAR $\uparrow$	FNIR $\downarrow$
<i>AdaFace*</i> [22]	♦	94.0	87.9	75.7	13.0 $\pm$ 3.5
CAL [11]	♠	81.4	74.7	66.3	52.8 $\pm$ 13.3
<i>BigGait*</i> [56]	♣	76.7	61.0	49.7	71.1 $\pm$ 6.1
<i>SapiensID</i> [25]	●	92.6	77.8	-	-
<i>Min-Fusion</i> [19]		87.1	79.2	62.4	48.5 $\pm$ 8.7
<i>Max-Fusion</i> [19]		89.9	89.3	73.4	23.0 $\pm$ 10.1
<i>Z-score</i> [46]		92.2	90.6	73.9	15.1 $\pm$ 1.5
<i>Min-max</i> [46]		91.8	90.9	73.9	15.4 $\pm$ 2.5
<i>Weighted-sum</i> [35]		91.7	90.6	73.6	15.4 $\pm$ 1.8
<i>Asym-AOI</i> [16]		92.3	90.0	74.0	15.9 $\pm$ 1.7
<i>BSSF</i> [50]	♦ ♠ ♣	91.8	91.1	73.9	14.1 $\pm$ 1.3
<i>FarSight</i> [28]		92.0	91.2	73.9	13.9 $\pm$ 1.1
<i>QME</i> [67]		<b>94.1</b>	90.8	76.2	12.3 $\pm$ 1.4
<i>FusionAgent (DA)</i>		92.8	<u>92.2</u>	<u>85.8</u>	<u>10.5 <math>\pm</math> 1.5</u>
<i>FusionAgent (CoT)</i>		<u>93.4</u>	<b>92.6</b>	<b>85.9</b>	<b>10.1 <math>\pm</math> 1.5</b>

Table 1. **Performance on CCVID.** [Keys: **Best** and second best performance; *Comb.*: model combination; \*: zero-shot performance; ♦: AdaFace for face modality; ♣: BigGait for gait modality; ♠: CAL of body modality; ●: SapiensID for face and body modalities; TAR: TAR@1%FAR; FNIR: FNIR@1%FPIR.]

<i>Method</i>	<i>Comb.</i>	Rank1 $\uparrow$	mAP $\uparrow$	TAR $\uparrow$	FNIR $\downarrow$
<i>AdaFace*</i> [22]	♦	18.5	5.9	2.4	99.8 $\pm$ 0.2
CAL [11]	♠	74.4	40.6	36.7	59.7 $\pm$ 7.3
AIM [54]	■	74.8	40.9	37.0	66.2 $\pm$ 7.5
<i>SapiensID</i> [25]	●	72.0	34.6	-	-
<i>Min-Fusion</i> [19]		38.1	13.5	12.4	81.9 $\pm$ 6.0
<i>Max-Fusion</i> [19]		62.5	33.3	16.8	94.8 $\pm$ 4.7
<i>Z-score</i> [46]		73.0	37.5	30.4	68.7 $\pm$ 9.2
<i>Min-max</i> [46]		73.2	38.1	31.9	75.1 $\pm$ 9.2
<i>Weighted-sum</i> [35]		73.2	37.8	31.3	72.4 $\pm$ 8.6
<i>Asym-AOI</i> [16]	♦ ♠ ■	71.2	32.9	19.1	76.3 $\pm$ 8.9
<i>BSSF</i> [50]		73.5	39.1	34.2	68.9 $\pm$ 8.5
<i>FarSight</i> [28]		73.2	37.8	31.3	72.4 $\pm$ 8.6
<i>QME</i> [67]		73.8	39.6	35.0	64.3 $\pm$ 8.0
<i>FusionAgent (DA)</i>		<b>75.5</b>	<b>41.0</b>	<u>36.5</u>	<u>50.3 <math>\pm</math> 9.0</u>
<i>FusionAgent (CoT)</i>		<b>75.5</b>	<b>41.0</b>	<b>37.0</b>	<b>50.0 <math>\pm</math> 8.5</b>

Table 2. **Performance on LTCC.** [Keys: **Best** and second best performance; *Comb.*: model combination; \*: zero-shot performance; ♦: AdaFace for face modality; ♠: CAL of body modality; ■: AIM for body modality; ●: SapiensID for face and body modalities; TAR: TAR@1%FAR; FNIR: FNIR@1%FPIR.]

for CCVID and  $k = 40$  for MEVID and LTCC. During training, we sample a continuous 4-frame clip from each video for CCVID and MEVID, and 1 frame for LTCC, as it is an image-based dataset. Training is conducted on 4 H100 GPUs with an effective batch size of 4, and takes nearly 4 hours. During evaluation, we disable sampling in FusionAgent to ensure results are reproducible.

#### 4.1. Experimental Results

**In-domain Evaluation.** As shown in Tab. 1, 2, and 3, the upper section of each table lists the unimodal models,

<i>Method</i>	<i>Comb.</i>	Rank1 $\uparrow$	mAP $\uparrow$	TAR $\uparrow$	FNIR $\downarrow$
<i>AdaFace*</i> [22]	♦	25.0	8.1	5.4	98.8 $\pm$ 1.2
CAL [11]	♠	52.5	27.1	34.7	67.8 $\pm$ 7.3
AGRL [52]	■	51.9	25.5	30.7	69.4 $\pm$ 8.9
<i>Min-Fusion</i> [19]		46.8	21.2	28.0	70.4 $\pm$ 8.0
<i>Max-Fusion</i> [19]		33.2	14.9	8.3	97.4 $\pm$ 1.6
<i>Z-score</i> [46]		54.1	27.4	30.7	66.5 $\pm$ 7.0
<i>Min-max</i> [46]		52.8	24.7	25.0	71.3 $\pm$ 6.1
<i>Weighted-sum</i> [35]		54.1	27.3	30.3	66.3 $\pm$ 7.0
<i>Asym-AOI</i> [16]	♦ ♠ ■	52.5	22.9	23.6	71.7 $\pm$ 5.8
<i>BSSF</i> [50]		53.5	27.4	30.5	65.9 $\pm$ 7.2
<i>FarSight</i> [28]		53.8	25.4	26.6	69.8 $\pm$ 6.4
<i>QME</i> [67]		<b>55.7</b>	<b>28.2</b>	<b>32.9</b>	<b>64.6 <math>\pm</math> 8.2</b>
<i>FusionAgent (DA)</i>		52.5	<b>28.7</b>	<u>34.8</u>	<u>60.8 <math>\pm</math> 7.3</u>
<i>FusionAgent (CoT)</i>		<u>54.7</u>	<b>28.7</b>	<b>34.9</b>	<b>58.6 <math>\pm</math> 7.4</b>

Table 3. **Performance on MEVID.** [Keys: **Best** and second best performance; *Comb.*: model combination; \*: zero-shot performance; ♦: AdaFace for face modality; ♠: CAL of body modality; ■: AGRL for body modality; TAR: TAR@1%FAR; FNIR: FNIR@1%FPIR.]

<i>Method</i>	Rank1 $\uparrow$	mAP $\uparrow$	TAR $\uparrow$	FNIR $\downarrow$
<i>CCVID <math>\rightarrow</math> LTCC (Zero-shot)</i>				
<i>FusionAgent (ACT)</i>	60.4	11.9	7.7	60.3 $\pm$ 8.6
<i>FusionAgent (FarSight)</i>	68.2	31.7	17.0	81.8 $\pm$ 9.6
<i>CCVID <math>\rightarrow</math> LTCC (10-shot)</i>				
<i>FusionAgent (ACT)</i>	73.6	39.8	34.8	53.5 $\pm$ 8.5
<i>MEVID <math>\rightarrow</math> LTCC (Zero-shot)</i>				
<i>FusionAgent (FarSight)</i>	75.3	42.3	37.7	59.3 $\pm$ 9.5
<i>FusionAgent (ACT)</i>	75.3	41.1	36.1	50.1 $\pm$ 8.4

Table 4. **Cross-domain Performance on LTCC.** FusionAgent is trained on CCVID with its model combination (Tab. 1) and evaluated on LTCC using LTCC’s model combination (Tab. 2). [Keys: TAR: TAR@1%FAR; FNIR: FNIR@1%FPIR.]

while the lower section reports the performance of fusion methods. Statistical baselines (*e.g.*, Z-score, FarSight) often fail to surpass the strongest single model across all metrics, while learning-based methods (*e.g.*, QME) achieve better gains and occasionally outperform the best unimodal model, but remain limited when fused models are not complementary (*e.g.*, LTCC). By contrast, FusionAgent consistently achieves superior performance on all three datasets, outperforming both unimodal and fusion baselines on most metrics. Across the tables, FNIR benefits the most from fusion, while Rank-1 improves the least. This is because FNIR is particularly sensitive to outliers, whereas Rank-1 is largely determined by the strongest modality and only marginally affected by fusion. Through top- $k$  selection and confidence weighting, our method effectively bounds the increase in non-match scores. On CCVID, the largest TAR gain arises because FusionAgent identifies high-quality facial inputs and primarily anchors on the FR model, which

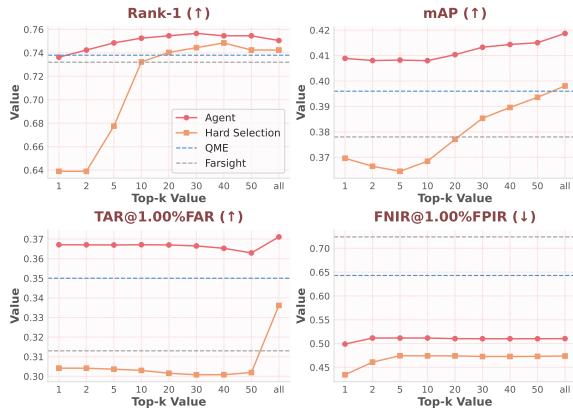


Figure 5. **Performance comparison on LTCC in four metrics.** FusionAgent consistently outperforms baselines, including the hard selection (*i.e.*, using all models), which highlights the effectiveness of dynamic model selection.

excels at distinguishing matches from non-matches. On LTCC, FNIR reduction is most pronounced due to the proposed ACT score-fusion strategy (see ablation in Tab. 5).

**Efficiency in Real-world Scenarios.** We introduce two inference modes: Direct Answering (DA) and Chain-of-Thought (CoT). DA avoids explicit reasoning, reducing average inference time from 2.81s to 1.03s on an H100 GPU, while maintaining competitive performance relative to QME (0.67s). CoT offers interpretable results, but at the cost of higher latency. Users can flexibly choose between the two to balance efficiency, interpretability, and accuracy.

**Cross-domain Evaluation.** In real-world deployments, agents may encounter unseen environments with different tool combinations, making cross-domain evaluation critical for assessing transferability. We evaluate FusionAgent by training on CCVID and testing on LTCC (Tab. 4). In the zero-shot setting, despite the changed tool pool, FusionAgent generalizes well and successfully executes tool calls without additional training. However, in CCVID  $\rightarrow$  LTCC, where face quality degrades, FusionAgent experiences a performance drop due to (1) the unknown robustness of newly introduced models or weights on the target dataset. Inappropriate anchor selection may amplify errors and degrade overall performance. (2) the domain gap in data distribution. With only 10 samples per subject (*i.e.*, 8% of the training data) and 50 training steps, FusionAgent quickly adapts and achieves performance comparable to full in-domain training. Under MEVID  $\rightarrow$  LTCC, the zero-shot setting attains similar performance, further demonstrating the robustness of FusionAgent.

## 4.2. Ablation Studies

**Agent Selection vs. Hard Selection.** We compare dynamic model selection by the agent (red line in Fig. 5) with hard

Method	Rank1	mAP	TAR@1%FAR	FNIR@1%FPiR
<i>QME</i> [67]	73.8	39.6	35.0	64.3 $\pm$ 8.0
<i>Z-score</i>	74.8	<b>41.7</b>	37.1	63.7 $\pm$ 9.5
<i>FarSight</i>	74.8	<b>41.7</b>	<b>37.2</b>	62.5 $\pm$ 9.7
<b><i>ACT (Ours)</i></b>	<b>75.5</b>	41.4	36.5	<b>51.0 <math>\pm</math> 9.4</b>

Table 5. **Comparison of score-fusion methods with agent selection on LTCC.** All fusion methods combined with agent-based selection outperform QME, confirming the value of dynamic model selection. The proposed ACT yields the best overall performance, particularly in open-set search.

selection (orange line), which uses all available models. Since hard selection lacks anchor information, we use the averaged scores as a surrogate anchor, whereas in the agent setting, the first selected model serves as the anchor. Our method yields notable improvements across all metrics over previous SoTA approaches, underscoring the efficacy of the proposed agent and the ACT algorithm. It also consistently outperforms hard selection in Rank-1, mAP, and TAR, highlighting the critical role of adaptive model selection. The significantly lower performance of hard selection relative to prior SoTA confirms that dynamic model selection is the key driver of performance gains.

**Top-K Values.** We investigate the impact of top- $k$  values in ACT. Fig. 5 indicates that the choice of  $k$  affects different metrics distinctively: larger  $k$  values generally improve Rank-1, mAP, and TAR@FAR, while FNIR remains relatively stable. Compared with hard selection, which exhibits large fluctuations, especially at small  $k$ , the proposed agent maintains both higher performance and greater stability across a wide range of  $k$  values. This shows that the agent not only adapts to sample-level variations by selectively exploiting complementary models but also avoids the noise introduced by averaging redundant scores. The robustness of ACT under varying  $k$  highlights its effectiveness in supporting adaptive model selection, ensuring consistent improvements across diverse evaluation settings.

**Alternative Score-fusion Methods.** Tab. 5 compares the proposed ACT against alternative statistical score-fusion methods using the same agent-selected model combinations. Even with standard fusion techniques such as Z-score and FarSight, our approach already surpasses the QME, highlighting the decisive role of dynamic model selection. Interestingly, Z-score and FarSight achieve nearly identical results, suggesting that the choice of statistical fusion has a limited impact once adaptive selection is applied. In contrast, the proposed ACT yields further consistent gains, with a substantial reduction in FNIR (down to 51.0). This demonstrates that beyond dynamic selection, robust score integration is crucial for handling challenging open-set search scenarios, making ACT more reliable and generalizable across diverse evaluation metrics.




Figure 6. **Statistics of model selection.** The first model is the anchor model for each category. The combination distribution reveals the characteristics of each dataset and which model combination provides robust performance.

Comb.	Rank1	mAP	TAR@1%FAR	FNIR@1%FPIR
<i>AIM+CAL</i>	74.8	41.0	36.1	53.4 ± 9.2
<i>FusionAgent</i>	75.5	41.4	36.5	51.0 ± 9.4

Table 6. **Ablation on model combination and anchor model selection on LTCC.** The results show that FusionAgent closely aligns with the grid-searched results but achieves slightly better performance due to sample-level diversity.




Figure 7. **Comparison of score distribution on CCVID.** [Keys: nm\_mean = mean value of non-match scores; mat\_mean = mean value of match scores.]

**Statistics of Dynamic Model Selection.** Fig. 6 presents the frequency of model and anchor selections per dataset. On CCVID, where faces are often clearly visible, the agent frequently selects AdaFace as the anchor and consistently combines it with CAL, while BigGait is selected less often, suggesting limited complementary value for combination. In contrast, LTCC and MEVID—collected under surveillance conditions with multi-view and low-quality faces—lead the agent to rely on ReID models.

**Dataset-level Model Selection.** Since exhaustive sample-level grid search is computationally prohibitive, we conduct grid search at the dataset level to determine the optimal model combination and anchor model. As shown in Tab. 6, AIM+CAL is the best model combination using ACT for the grid search, which is highly consistent with the combination chosen by our agent. Nevertheless, the agent achieves superior performance, attributed to its ability to exploit sample-level diversity during model selection.

### 4.3. Qualitative Results

**Score Distribution.** Fig. 7 plots the distribution of non-match (imposter) scores, match (genuine) scores, and the threshold for FAR=1% for QME and FusionAgent. For

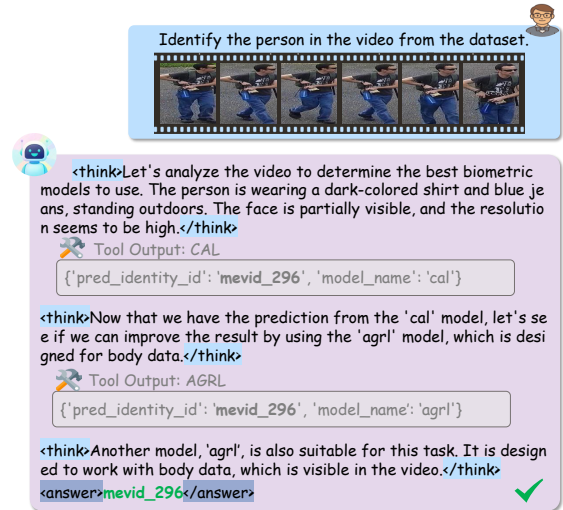


Figure 8. **CoT of FusionAgent on MEVID.** FusionAgent performs interpretable reasoning and dynamically selects the suitable tool combination for each test sample.

clarity, we align the y-axis and normalize the scores. Our method demonstrates a markedly larger margin between the match scores and the FAR threshold, while effectively suppressing non-match scores near zero. This improvement stems from two key mechanisms: (1) by summing only the top- $k$  confident scores, we amplify the contribution of true matches and widen the separation from non-matches, which are predominantly excluded; and (2) our dynamic model selection provides a more robust set of models per query, leading to higher-quality scores for fusion.

**CoT for Model Selection.** Fig. 8 illustrates the CoT mode of FusionAgent, where the agent interprets the input and dynamically selects suitable models for each query sample. By leveraging the learned reliability and complementarity of different models, FusionAgent reasons over the query’s characteristics to determine the optimal model combination. This CoT process enhances interpretability by explicitly revealing the agent’s decision path and selection rationale.

## 5. Conclusion

We propose FusionAgent, a novel agentic framework for human recognition with dynamic model selection. We introduce an Anchor-based Confidence Top- $k$  score-fusion method (ACT) for sample-dependent and adaptive model integration. Multiple reward functions are designed to guide the agent’s tool usage and exploration of diverse model combinations. Extensive experiments and analyses validate the effectiveness of FusionAgent and ACT, highlighting the benefit of query-wise model selection and fusion. Our approach provides a scalable and extensible solution for multi-modal and multi-model tasks, with potential applicability to broader vision tasks [10, 12, 26, 31, 43, 60, 65, 66]. We hope this work inspires further research on agentic AI.

## Acknowledgments

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2022-21102100004. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

## References

- [1] John Arevalo, Tamar Solorio, Manuel Montes-y Gómez, and Fabio A González. Gated multimodal units for information fusion. *arXiv preprint arXiv:1702.01992*, 2017. 2
- [2] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 5
- [3] Boyu Chen, Zhengrong Yue, Siran Chen, Zikang Wang, Yang Liu, Peng Li, and Yali Wang. Lvagent: Long video understanding by multi-round dynamical collaboration of mllm agents. *arXiv preprint arXiv:2503.10200*, 2025. 2, 3
- [4] Junwen Chen, Jie Zhu, and Yu Kong. Atm: Action temporality modeling for video question answering. In *ACM MM*, 2023. 1
- [5] Mohamed Cheniti, Zahid Akhtar, Chandranath Adak, and Kamran Siddique. An approach for full reinforcement-based biometric score fusion. *IEEE Access*, 2024. 2
- [6] Daniel Davila, Dawei Du, Bryon Lewis, Christopher Funk, Joseph Van Pelt, Roderic Collins, Kellie Corona, Matt Brown, Scott McCloskey, Anthony Hoogs, et al. Mevid: Multi-view extended videos with identities for video person re-identification. In *WACV*, 2023. 5
- [7] Maria De Marsico, Michele Nappi, and Daniel Riccio. Cabala—collaborative architectures based on biometric adaptable layers and activities. *PR*, 2012. 2
- [8] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 2
- [9] Jiankang Deng, Jia Guo, Tongliang Liu, Mingming Gong, and Stefanos Zafeiriou. Sub-center arcface: Boosting face recognition by large-scale noisy web faces. In *ECCV*, 2020. 2
- [10] Xingyu Fu, Siyi Liu, Yinuo Xu, Pan Lu, Guangqiuse Hu, Tianbo Yang, Taran Anantasagar, Christopher Shen, Yikai Mao, Yuanzhe Liu, et al. Learning human-perceived fakeness in ai-generated videos via multimodal llms. *arXiv preprint arXiv:2509.22646*, 2025. 8
- [11] Xinqian Gu, Hong Chang, Bingpeng Ma, Shutao Bai, Shiguang Shan, and Xilin Chen. Clothes-changing person re-identification with rgb modality only. In *CVPR*, 2022. 2, 5, 6
- [12] Xiao Guo, Xiufeng Song, Yue Zhang, Xiaohong Liu, and Xiaoming Liu. Rethinking vision-language model in face forensics: Multi-modal interpretable forged face detector. In *CVPR*, 2025. 1, 8
- [13] Xiao Guo, Yue Zhang, Mohit Bansal, and Xiaoming Liu. Deepfake-agent: Aggregating semantic forgery clues for generalizable detection. In *In Proceeding of Findings of IEEE Computer Vision and Pattern Recognition*, 2026. 2, 3
- [14] Xiao Guo, Jie Zhu, Anil Jain, and Xiaoming Liu. On the holistic approach for detecting human image forgery. *arXiv preprint arXiv:2601.04715*, 2026. 1
- [15] Mingxing He, Shi-Jinn Horng, Pingzhi Fan, Ray-Shine Run, Rong-Jian Chen, Jui-Lin Lai, Muhammad Khurram Khan, and Kevin Octavius Sentosa. Performance evaluation of score level fusion in multimodal biometric systems. *PR*, 2010. 2, 5
- [16] Abderrahmane Herbadji, Zahid Akhtar, Kamran Siddique, Noubel Guermat, Lahcene Ziet, Mohamed Cheniti, and Khan Muhammad. Combining multiple biometric traits using asymmetric aggregation operators for improved person recognition. *Symmetry*. 5, 6
- [17] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 2022. 5
- [18] Zhanbo Huang, Dingqiang Ye, Xiaoming Liu, and Yu Kong. Unlocking motion from large vision models with a semantic and kinematic duality for gait recognition. In *CVPR*, 2026. 1, 2
- [19] Anil Jain, Karthik Nandakumar, and Arun Ross. Score normalization in multimodal biometric systems. *PR*, 2005. 1, 2, 5, 6
- [20] Anil Jain, Ruud Bolle, and Sharath Pankanti. Introduction to biometrics. In *Biometrics: personal identification in networked society*. Springer, 2011. 2
- [21] Yuyang Ji, Yixuan Shen, Kien Nguyen, Lifeng Zhou, and Feng Liu. Idselect: A rl-based cost-aware selection agent for video-based multi-modal person recognition. *arXiv preprint arXiv:2602.18990*, 2026. 2
- [22] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *CVPR*, 2022. 2, 6
- [23] Minchul Kim, Feng Liu, Anil K Jain, and Xiaoming Liu. Cluster and aggregate: Face recognition with large probe set. In *NeurIPS*, 2022.
- [24] Minchul Kim, Yiyang Su, Feng Liu, Anil Jain, and Xiaoming Liu. KeyPoint Relative Position Encoding for Face Recognition. In *CVPR*, 2024. 2
- [25] Minchul Kim, Dingqiang Ye, Yiyang Su, Feng Liu, and Xiaoming Liu. Sapiensid: Foundation for human recognition. In *CVPR*, 2025. 1, 2, 5, 6
- [26] Abhinav Kumar, Yuliang Guo, Zhihao Zhang, Xinyu Huang, Liu Ren, and Xiaoming Liu. Charm3r: Towards unseen camera height robust monocular 3d detector. In *CVPR*, 2025. 1, 8
- [27] Feng Liu, Minchul Kim, ZiAng Gu, Anil Jain, and Xiaoming Liu. Learning clothing and pose invariant 3d shape representation for long-term person re-identification. In *ICCV*, 2023. 2

- [28] Feng Liu, Ryan Ashbaugh, Nicholas Chimitt, Najmul Hassan, Ali Hassani, Ajay Jaiswal, Minchul Kim, Zhiyuan Mao, Christopher Perry, Zhiyuan Ren, et al. Farsight: A physics-driven whole-body biometric system at large distance and altitude. In *WACV*, 2024. 2, 5, 6
- [29] Feng Liu, Minchul Kim, Zhiyuan Ren, and Xiaoming Liu. Distilling CLIP with Dual Guidance for Learning Discriminative Human Body Shape Representation. In *CVPR*, 2024. 2
- [30] Feng Liu, Nicholas Chimitt, Lanqing Guo, Jitesh Jain, Aditya Kane, Minchul Kim, Wes Robbins, Yiyang Su, Dingqiang Ye, Xingguang Zhang, et al. Person recognition at altitude and range: Fusion of face, body shape and gait. *arXiv preprint arXiv:2505.04616*, 2025. 2
- [31] Yuanzhe Liu, Jingyuan Zhu, Yuchen Mo, Gen Li, Xu Cao, Jin Jin, Yifan Shen, Zhengyuan Li, Tianjiao Yu, Wenzhen Yuan, et al. Palm: Progress-aware policy learning via affordance reasoning for long-horizon robotic manipulation. *arXiv preprint arXiv:2601.07060*, 2026. 8
- [32] Ziyu Liu, Zeyi Sun, Yuhang Zang, Xiaoyi Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi Wang. Visual-rlft: Visual reinforcement fine-tuning. *arXiv preprint arXiv:2503.01785*, 2025. 3, 4
- [33] Ziyu Liu, Yuhang Zang, Yushan Zou, Zijian Liang, Xiaoyi Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi Wang. Visual agentic reinforcement fine-tuning. *arXiv preprint arXiv:2505.14246*, 2025. 2, 3, 4
- [34] Karthik Nandakumar, Yi Chen, Sarat C Dass, and Anil Jain. Likelihood ratio-based biometric score fusion. *TPAMI*, 2007. 2
- [35] Tae Jin Park, Manoj Kumar, and Shrikanth Narayanan. Multi-scale speaker diarization with neural affinity score fusion. In *ICASSP*, 2021. 1, 2, 5, 6
- [36] Juan-Manuel Pérez-Rúa, Valentin Vielzeuf, Stéphane Pateux, Moez Baccouche, and Frédéric Jurie. Mfas: Multimodal fusion architecture search. In *CVPR*, 2019. 2
- [37] Norman Poh and Josef Kittler. A unified framework for biometric expert fusion incorporating quality measures. *TPAMI*, 2011. 2
- [38] Norman Poh, Josef Kittler, and Thirimachos Bourlai. Improving biometric device interoperability by likelihood ratio-based quality dependent score normalization. In *BTAS*, 2007. 2
- [39] Xuelin Qian, Wenxuan Wang, Li Zhang, Fangrui Zhu, Yanwei Fu, Tao Xiang, Yu-Gang Jiang, and Xiangyang Xue. Long-term cloth-changing person re-identification. In *ACCV*, 2020. 5
- [40] Arun Ross and Anil Jain. Information fusion in biometrics. *PR letters*, 2003. 2
- [41] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024. 3, 4, 5
- [42] Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun Zhang, Kangjia Zhao, Qianqian Zhang, et al. Vlm-r1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*, 2025. 3
- [43] Yifan Shen, Yuanzhe Liu, Jingyuan Zhu, Xu Cao, Xiaofeng Zhang, Yixiao He, Wenming Ye, James Matthew Rehg, and Ismeni Lourentzou. Fine-grained preference optimization improves spatial reasoning in vlms. *arXiv preprint arXiv:2506.21656*, 2025. 8
- [44] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *NeurIPS*, 2023. 3
- [45] Maneet Singh, Richa Singh, and Arun Ross. A comprehensive overview of biometric fusion. *Information Fusion*, 52, 2019. 2
- [46] Robert Snelick, Mike Indovina, James Yen, and Alan Mink. Multimodal biometrics: issues in design and testing. In *ICMI*, 2003. 1, 5, 6
- [47] Yiyang Su, Minchul Kim, Feng Liu, Anil Jain, and Xiaoming Liu. Open-set biometrics: Beyond good closed-set models. In *ECCV*, 2024. 5
- [48] Yiyang Su, Yunping Shi, Feng Liu, and Xiaoming Liu. Hamobe: Hierarchical and adaptive mixture of biometric experts for video-based person reid. In *ICCV*, 2025. 2
- [49] Yiyang Su, Minchul Kim, Jie Zhu, Christopher Perry, Feng Liu, Anil Jain, and Xiaoming Liu. Localscore: Local density-aware similarity scoring for biometrics. *arXiv preprint arXiv:2602.01012*, 2026. 2
- [50] Jackson Horlick Teng, Thian Song Ong, Tee Connie, Kalaiarasi Sonai Muthu Anbananthen, and Pa Pa Min. Optimized score level fusion for multi-instance finger vein recognition. *Algorithms*, 2022. 1, 2, 5, 6
- [51] Mayank Vatsa, Richa Singh, and Afzel Noore. Integrating image quality in 2v-svm biometric match score fusion. *International Journal of Neural Systems*, 2007. 2
- [52] Yiming Wu, Omar El Farouk Bourahla, Xi Li, Fei Wu, Qi Tian, and Xue Zhou. Adaptive graph representation learning for video person re-identification. *TIP*, 2020. 2, 6
- [53] Feng Xu, Guangyao Zhai, Xin Kong, Tingzhong Fu, Daniel FN Gordon, Xueli An, and Benjamin Busam. Stare-vla: Progressive stage-aware reinforcement for fine-tuning vision-language-action models. *arXiv preprint arXiv:2512.05107*, 2025. 3
- [54] Zhengwei Yang, Meng Lin, Xian Zhong, Yu Wu, and Zheng Wang. Good is bad: Causality inspired cloth-debiasing for cloth-changing person re-identification. In *CVPR*, 2023. 2, 6
- [55] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *ICLR*, 2023. 2, 3, 4
- [56] Dingqiang Ye, Chao Fan, Jingzhe Ma, Xiaoming Liu, and Shiqi Yu. BigGait: Learning Gait Representation You Want by Large Vision Models. In *CVPR*, 2024. 2, 6
- [57] Dingqiang Ye, Chao Fan, Zhanbo Huang, Chengwen Luo, Jianqiang Li, Shiqi Yu, and Xiaoming Liu. Biggergait: Unlocking gait recognition with layer-wise representations from large vision models. In *NeurIPS*, 2025. 2

- [58] Mustafa Berkay Yılmaz and Berrin Yanıkoğlu. Score level fusion of classifiers in off-line signature verification. *Information Fusion*, 2016. 2
- [59] En Yu, Kangheng Lin, Liang Zhao, Jisheng Yin, Yana Wei, Yuang Peng, Haoran Wei, Jianjian Sun, Chunrui Han, Zheng Ge, et al. Perception-r1: Pioneering perception policy with reinforcement learning. *arXiv preprint arXiv:2504.07954*, 2025. 3
- [60] Yuan Yuan, Zhaojian Li, and Bin Zhao. A survey of multimodal learning: Methods, applications, and future. *ACM Computing Surveys*, 57(7):1–34, 2025. 8
- [61] Ziyuan Zhang, Luan Tran, Xi Yin, Yousef Atoum, Xiaoming Liu, Jian Wan, and Nanxin Wang. Gait recognition via disentangled representation learning. In *CVPR*, 2019. 2
- [62] Ziyuan Zhang, Luan Tran, Feng Liu, and Xiaoming Liu. On learning disentangled representations for gait recognition. *TPAMI*, 2020. 2
- [63] Zhihao Zhang, Yiwei Chen, Weizhan Zhang, Caixia Yan, Qinghua Zheng, Qi Wang, and Wangdu Chen. Tile classification based viewport prediction with multi-modal fusion transformer. In *ACM MM*, 2023. 1, 2
- [64] Zhihao Zhang, Shengcao Cao, and Yu-Xiong Wang. Tamm: Triadapter multi-modal learning for 3d shape understanding. In *CVPR*, 2024. 1
- [65] Zhihao Zhang, Abhinav Kumar, Girish Chandar Ganesan, and Xiaoming Liu. Unleashing the power of chain-of-prediction for monocular 3d object detection. In *CVPR*, 2026. 1, 8
- [66] Zhihao Zhang, Abhinav Kumar, and Xiaoming Liu. Towards intrinsic-aware monocular 3d object detection. In *CVPR*, 2026. 1, 8
- [67] Jie Zhu, Yiyang Su, Minchul Kim, Anil Jain, and Xiaoming Liu. A quality-guided mixture of score-fusion experts framework for human recognition. In *ICCV*, 2025. 1, 2, 5, 6, 7
- [68] Jie Zhu, Yiyang Su, and Xiaoming Liu. Can textual reasoning improve the performance of mllms on fine-grained visual classification? *arXiv preprint arXiv:2601.06993*, 2026. 3, 4
- [69] Yushen Zuo, Qi Zheng, Mingyang Wu, Xinrui Jiang, Renjie Li, Jian Wang, Yide Zhang, Gengchen Mai, Lihong V Wang, James Zou, et al. 4kagent: agentic any image to 4k super-resolution. *arXiv preprint arXiv:2507.07105*, 2025. 2, 3