

The Good, the Better, and the Best: Improving the Discriminability of Face Embeddings through Attribute-aware Learning

Ana Dias^{1,2,3}, João Ribeiro Pinto⁴, Hugo Proença^{1,2}, João C. Neves^{1,3}

¹University of Beira Interior, Portugal ²IT: Instituto de Telecomunicações ³NOVA LINCS ⁴Amadeus, Portugal
ana.margarida.dias@ubi.pt

Abstract—Despite recent advances in face recognition, robust performance remains challenging under large variations in age, pose, and occlusion. A common strategy to address these issues is to guide representation learning with auxiliary supervision from facial attributes, encouraging the visual encoder to focus on identity-relevant regions. However, existing approaches typically rely on heterogeneous and fixed sets of attributes, implicitly assuming equal relevance across attributes. This assumption is suboptimal, as different attributes exhibit varying discriminative power for identity recognition, and some may even introduce harmful biases. In this paper, we propose an attribute-aware face recognition architecture that supervises the learning of facial embeddings using identity class labels, identity-relevant facial attributes, and non-identity-related attributes. Facial attributes are organized into interpretable groups, making it possible to decompose and analyze their individual contributions in a human-understandable manner. Experiments on standard face verification benchmarks demonstrate that joint learning of identity and facial attributes improves the discriminability of face embeddings with two major conclusions: (i) using identity-relevant subsets of facial attributes consistently outperforms supervision with a broader attribute set, and (ii) explicitly forcing embeddings to *unlearn* non-identity-related attributes yields further performance gains compared to leaving such attributes unsupervised. Additionally, our method serves as a diagnostic tool for assessing the trustworthiness of face recognition encoders by allowing for the measurement of accuracy gains with suppression of non-identity-relevant attributes, with such gains suggesting shortcut learning from redundant attributes associated with each identity.

Index Terms—Face Recognition, Multi-Task Learning, Facial Attribute Supervision, Attribute Suppression

I. INTRODUCTION

Despite significant advances over the past decade [1]–[3], face recognition methods still struggle to correctly identify individuals under factors such as age variation [4], [5], pose changes [6], [7], and occlusion [8], [9]. To address these limitations, a common strategy is the use of auxiliary supervision during training, usually with facial attributes, to guide representation learning toward identity-related cues. This is typically achieved within a multi-task learning framework, in which a shared visual encoder is jointly optimized for identity recognition and facial attribute estimation [10]–[13]. However, existing multi-task face recognition approaches usually treat facial attributes as a single, homogeneous auxiliary objective defined over a fixed attribute set. This formulation neglects




Fig. 1. **From fixed attribute supervision to selective prediction and suppression.** (a) Prior multi-task face recognition methods use a fixed attribute list as an auxiliary prediction task alongside identity learning. (b) Our approach organizes attributes into groups and controls their influence during training: groups can be predicted or selectively suppressed, while the identity learning objective remains unchanged.

two key aspects: (a) facial attributes differ substantially in their relevance and discriminative power for identity recognition, and some may even introduce noise or bias into the learned representation; and (b) actively suppressing certain attributes can be more beneficial than simply leaving them unsupervised. Motivated by these observations, we introduce a face recognition architecture that enhances the discriminability of identity-related features by encouraging the encoder to preserve information associated with identity-relevant facial attributes while ignoring or suppressing information related to attributes that are irrelevant for recognition (Figure 1). To this end, we adopt a state-of-the-art visual encoder supervised

by three complementary loss signals: identity recognition, identity-relevant facial attributes, and identity-irrelevant facial attributes. The first two loss signals are jointly optimized using standard gradient descent to promote the learning of identity-discriminative features. In contrast, the loss associated with identity-irrelevant attributes is combined with a gradient reversal mechanism, encouraging the model to explicitly “unlearn” attribute information that is detrimental to identity recognition. Furthermore, facial attributes are organized into interpretable, region-based groups, enabling a systematic and human-understandable analysis of how different attribute subsets contribute to identity learning. Using the proposed architecture, we demonstrate that jointly learning identity and facial attributes is clearly beneficial to identity-only supervision, but, as hypothesized by us, this increment is suboptimal, since identity-relevant attribute groups yields stronger and more stable improvements than using a heterogeneous mix of attributes. Moreover, we show that explicitly suppressing non-identity-related attributes provides even further performance gains, supporting the hypothesis that actively suppressing weakly relevant attributes can be more effective than simply leaving them unsupervised. While most attribute groups support these conclusions, attributes related to *hair* deviate from the general trend. Despite being non-identity-relevant, its suppression degrades recognition performance, indicating that this region encodes identity-discriminative information. We hypothesize that this is due to the low intra-subject variability of hair-related attributes in common face datasets, and we demonstrate this claim by achieving $\approx 80\%$ verification accuracy solely by using hair-related attributes, without any visual input. These findings expose attribute–identity correlations that are often overlooked in face recognition benchmarks, and demonstrate how our method can be leveraged to identify and analyze such biases. Accordingly, beyond improving recognition performance, the proposed approach provides a useful tool for promoting more transparent and trustworthy face recognition systems by explicitly revealing which attributes contribute to identity discrimination.

In summary, our main contributions are as follows:

- We introduce an attribute-aware face recognition architecture that jointly learns identity representations and facial attributes, while explicitly suppressing non-identity-related attribute information via a gradient reversal mechanism.
- We experimentally demonstrate that supervising selected identity-relevant attribute groups, together with explicit suppression of non-identity attributes, consistently outperforms joint learning identity with a heterogeneous mix of facial attributes.
- We show that selective attribute suppression not only improves recognition performance, but also helps expose shortcut learning in attribute-supervised face recognition, evidencing the potential of our approach for analyzing dataset biases and promoting more transparent and trustworthy recognition systems.

II. STATE OF ART

A. Multi-task Face Recognition

Early research on multi-task face recognition explored the use of facial attributes to improve identity learning through auxiliary supervision. Wang et al. [11] proposed a joint optimization framework for face recognition and attribute prediction and reported improved recognition performance, while Taherkhani et al. [10] showed that attributes provide complementary information for identity recognition. Zhong et al. [14] further showed that face recognition representations capture attribute-related information, consistent with the prior benefits of attribute supervision. More recent frameworks extend this idea to multi-task face perception models with shared backbones and task-specific heads, jointly training face recognition with tasks such as attribute prediction, age estimation, and expression recognition [12], [13]. While these methods demonstrate the effectiveness of joint learning, they typically adopt a fixed attribute list and treat attribute supervision as a single auxiliary objective. Consequently, the effects of different attribute groups, and their selective prediction or suppression, on face recognition performance remain underexplored. Our multitask learning architecture enables controlled prediction and suppression of selected attribute groups and explores their effects on face recognition performance.

B. Adversarial Learning for Invariance

Gradient Reversal Layers (GRLs) [15] have been adopted in multi-task face recognition to suppress attribute-specific information, most commonly age, from identity representations by reversing the gradient from an auxiliary branch during backpropagation. MTLFace [4] integrates a recognition discriminator with a GRL [15] to encourage the learned identity feature to be age-invariant, while jointly optimizing age estimation and face recognition objectives. OrdCon [16] follows a similar strategy and gradually applies a GRL to an age branch to enforce age-invariant identity features. More generally, this same mechanism can be applied to promote invariance to additional facial attributes by attaching an attribute classifier to the identity embedding and training it through a GRL. In our work, we adopt this approach and extend GRL-based suppression beyond age to multiple facial attribute groups.

III. METHODOLOGY

This section presents the proposed attribute-aware training architecture for face recognition. We first describe the overall model architecture, followed by the mechanisms used for predicting and suppressing attributes. An overview of the proposed architecture is illustrated in Figure 2.

A. Network Architecture

Given an input face image x , a shared feature extractor $f_\theta(\cdot)$ produces an identity embedding

$$\mathbf{z} = f_\theta(x). \quad (1)$$

The embedding $\mathbf{z}$ is subsequently processed by an identity classification head $h_{\text{id}}(\cdot)$ and a set of group-specific attribute




Fig. 2. **Overview of the proposed attribute-aware face recognition architecture.** A shared visual encoder produces an identity embedding used for identity recognition and for attribute-based auxiliary supervision. Attribute groups selected for prediction are supervised directly, while groups selected for suppression are connected through gradient reversal layers (GRL) to discourage attribute-specific information in the identity embedding.

heads $\{h_g(\cdot)\}$, one for each attribute group g . Attribute heads operate in two modes: (i) prediction, where h_g receives $\mathbf{z}$ directly, and (ii) suppression, where $\mathbf{z}$ is first passed through a gradient reversal layer (GRL).

B. Attribute Group Prediction

We select a set of attribute groups to be trained in prediction mode, denoted as:

$$\mathcal{G}_{\text{pred}} = \{G_1, \dots, G_K\}, \quad (2)$$

where each group G_k is a set of attributes. For each group $G_k \in \mathcal{G}_{\text{pred}}$, the corresponding attribute head produces group-specific logits:

$$\hat{\mathbf{a}}_k = h_k(\mathbf{z}), \quad (3)$$

where $\hat{\mathbf{a}}_k \in \mathbb{R}^{|G_k|}$ contains one logit per attribute in group G_k . The ground-truth binary label of the i^{th} attribute in G_k is denoted by $y_{k,i} \in \{0, 1\}$ and the corresponding predicted logit as $\hat{a}_{k,i}$.

Using the binary cross-entropy (BCE) loss, the attribute prediction loss is defined as:

$$\mathcal{L}_{\text{attr}} = \frac{1}{\sum_{k=1}^K |G_k|} \sum_{k=1}^K \sum_{i \in G_k} \text{BCE}(\hat{a}_{k,i}, y_{k,i}), \quad (4)$$

where $|G_k|$ denotes the number of attributes in group G_k .

C. Attribute Group Suppression

We select a set of attribute groups for suppression, disjoint from those used for prediction, denoted as:

$$\mathcal{G}_{\text{sup}} = \{G_1, \dots, G_M\}, \quad (5)$$

where each group G_m is a set of attributes. For each group $G_m \in \mathcal{G}_{\text{sup}}$, the corresponding attribute head receives the identity embedding through a gradient reversal layer (GRL):

$$\hat{\mathbf{a}}_m^{\text{adv}} = h_m^{\text{adv}}(\text{GRL}(\mathbf{z})), \quad (6)$$

where $\hat{\mathbf{a}}_m^{\text{adv}} \in \mathbb{R}^{|G_m|}$ contains one logit per attribute in group G_m . For each attribute indexed by $i \in G_m$, we denote its

ground-truth binary label as $y_{m,i} \in \{0, 1\}$ and the corresponding predicted logit as $\hat{a}_{m,i}^{\text{adv}}$. The adversarial suppression objective is defined as:

$$\mathcal{L}_{\text{adv}} = \frac{1}{\sum_{m=1}^M |G_m|} \sum_{m=1}^M \sum_{i \in G_m} \text{BCE}(\hat{a}_{m,i}^{\text{adv}}, y_{m,i}), \quad (7)$$

where $|G_m|$ denotes the number of attributes in group G_m .

During backpropagation, the GRL reverses gradients from $\mathcal{L}_{\text{adv}}$ to $\mathbf{z}$, thereby encouraging the identity embedding to discard information related to the suppressed attribute groups.

D. Multi-task Learning Framework

We optimize face recognition jointly with the attribute objectives introduced above. The identity classification loss serves as the primary objective, with attribute prediction and suppression losses incorporated as specified below. Identity recognition is supervised using the CosFace loss [1], defined as:

$$\mathcal{L}_{\text{id}} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M \log \frac{e^{r(\cos(\theta_{y_i,i})-m)}}{e^{r(\cos(\theta_{y_i,i})-m)} + \sum_{j \neq y_i} e^{r \cos(\theta_{j,i})}}, \quad (8)$$

where N is the batch size, M is the number of identities, $\theta_{j,i}$ denotes the angle between the normalized feature of sample x_i and the normalized prototype of identity j , m is the angular margin, and r is the scale factor.

For attribute groups trained in prediction mode, the corresponding heads are optimized using $\mathcal{L}_{\text{attr}}$. For attribute groups trained in suppression mode, we additionally optimize the adversarial loss $\mathcal{L}_{\text{adv}}$.

The overall training objective is given by:

$$\mathcal{L} = \mathcal{L}_{\text{id}} + \lambda_{\text{pred}} \mathcal{L}_{\text{attr}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}}, \quad (9)$$

where $\mathcal{L}_{\text{id}}$ denotes the identity loss, $\mathcal{L}_{\text{attr}}$ the attribute prediction loss, and $\mathcal{L}_{\text{adv}}$ the adversarial loss used for attribute suppression. The weights λ_{pred} and λ_{adv} control the strength of attribute prediction and suppression, respectively.

IV. EXPERIMENTAL SETUP

A. Attribute Selection, Grouping, and Experimental Design

We conduct all experiments using the MAAD-Face dataset [17], which provides large-scale identity annotations together with binary facial attribute labels. Compared to datasets such as CelebA and LFW [18], MAAD-Face offers greater scale and more reliable identity annotations, making it suitable for analyzing the effect of attribute-based supervision on face recognition performance.

We organize attribute supervision into region-based groups to study how supervision from different parts of the face contributes to identity recognition, while keeping the analysis focused on interpretable facial areas rather than individual attribute types. Our goal is to contrast sources of supervision that differ in their expected relevance to identity learning. We therefore define five region-based groups, namely PERIOULAR, MOUTH, NOSE, HAIR, and ACCESSORIES. This grouping includes attributes associated with relatively stable facial regions (PERIOULAR, MOUTH, NOSE) alongside attributes that are more variable or potentially misleading (ACCESSORIES), as well as attributes that may introduce dataset-specific identity correlations (HAIR). From the 47 attributes available in MAAD-Face, we retain only those that map directly to these groups, resulting in 19 attributes, as detailed in Table I.

TABLE I
REGION-BASED ATTRIBUTE GROUPS. ATTRIBUTE SUPERVISION IS ORGANIZED INTO FIVE REGION-BASED GROUPS, NAMELY PERIOULAR (P), MOUTH (M), NOSE (N), HAIR (H), AND ACCESSORIES (A). THE ATTRIBUTES INCLUDED IN EACH GROUP ARE LISTED.

Groups	Attributes
Periocular (P)	Bags Under Eyes, Bushy Eyebrows, Arched Eyebrows, No Eyewear, Eyeglasses
Mouth (M)	Big Lips, Wearing Lipstick
Nose (N)	Big Nose, Pointy Nose
Hair (H)	Bald, Wavy Hair, Receding Hairline, Bangs, Sideburns, Black Hair, Blond Hair, Gray Hair
Accessories (A)	Wearing Hat, Wearing Earrings

Using the framework described in Section III, we consider three experimental settings:

- 1) *Single-group prediction.* We jointly train face recognition and attribute prediction for one group at a time, allowing us to isolate the contribution of each attribute group to recognition performance.
- 2) *Multiple-group prediction.* We jointly train face recognition and attribute prediction across multiple groups simultaneously, allowing us to study interactions between different attribute groups and their combined effect on performance.
- 3) *Multiple-group prediction with suppression.* We jointly train face recognition with attribute prediction across multiple groups while suppressing selected ones, allowing us to evaluate how reducing information from specific attributes affects recognition performance.

B. Implementation Details

Training Data. All models are trained on MAAD-Face [17]. Since some attributes contain missing annotations, we restrict training to the largest subset of images for which all selected attributes are defined. This results in approximately 252k training images, which are shared across all experiments.

Preprocessing. For both training and evaluation images, faces are detected and aligned using RetinaFace [19], with MTCNN [20] used as a fallback. Aligned faces are cropped and resized to a fixed resolution of 112×112 pixels.

Face Verification Benchmarks. Face recognition performance is evaluated on four standard verification benchmarks: CALFW [21], AgeDB [22], CACD-VS [23], and ECAF [4]. All benchmarks follow standard 10-fold cross-validation.

Training Details. We use a Swin Transformer [24], a state-of-the-art vision encoder typically used in face recognition [12], as the shared feature extractor in all experiments. All models are trained using stochastic gradient descent (SGD) with an initial learning rate of 0.01, reduced by a factor of 10 at epochs 20 and 35. We use a batch size of 512 across all experiments. The CosFace [1] loss is configured with margin $m = 0.35$ and scale $s = 64.0$. We set $\lambda_{\text{pred}} = 5$ and $\lambda_{\text{adv}} = 2$ based on preliminary validation experiments. Early stopping is applied with a patience of 20 epochs based on validation performance. All experiments are conducted on an NVIDIA A40 GPU.

V. RESULTS

A. Effects of Attribute-Based Supervision on Face Recognition

We evaluate the proposed framework across the three experimental settings defined in Section IV, using verification accuracy on the CALFW [21], AgeDB [22], CACD-VS [23], and ECAF [4] benchmarks, and report the average across them. For ECAF [4], we report two protocols: $\langle \text{Adult}, \text{Child} \rangle$ (ECAF_{a2c}) and $\langle \text{Child}, \text{Child} \rangle$ (ECAF_{c2c}) verification. All configurations are compared against a face recognition baseline trained without attribute supervision, which allows us to isolate the effects of attribute prediction and suppression.

Single-group prediction. Table II reports results from jointly training face recognition and a single attribute group at a time. This setting isolates each group’s individual contribution to recognition performance. Overall, attribute groups associated with relatively stable facial regions, particularly PERIOULAR and NOSE, consistently improve performance over the baseline across most benchmarks, indicating that this supervision shifts the model’s focus toward regions that are reliably informative of identity. MOUTH provides smaller but generally positive gains, potentially due to higher variability from expression or cosmetics. ACCESSORIES provides limited and inconsistent improvements, consistent with the expectation that such attributes are weakly related to identity and may introduce noise rather than complementary information. Surprisingly, HAIR also improves performance on several benchmarks, even though hair-related appearance is more variable and less directly tied to identity than core facial regions. This pattern

may reflect dataset-specific correlations between hair attributes and identity, an effect we analyze further in later sections.

TABLE II
VERIFICATION ACCURACY (%) WITH SINGLE-GROUP ATTRIBUTE SUPERVISION. FR DENOTES THE FACE RECOGNITION BASELINE TRAINED WITHOUT ATTRIBUTE SUPERVISION. “+P/+M/+N/+H/+A” INDICATES JOINT TRAINING WITH PREDICTION OF ONE ATTRIBUTE GROUP (PERIOCLAR/MOUTH/NOSE/HAIR/ACCESSORIES). AVG. CORRESPONDS TO THE AVERAGE VERIFICATION ACCURACY ACROSS BENCHMARKS.

Mode	Avg.	ECAF _{c2c}	ECAF _{a2c}	AgeDB	CALFW	CACD-VS
FR	84.13	79.35	75.08	81.97	86.58	<u>97.65</u>
+P	<u>84.65</u>	<u>79.75</u>	<u>76.05</u>	83.57	86.22	97.68
+M	84.39	79.60	75.98	83.12	85.68	97.58
+N	84.70	80.20	75.93	83.33	<u>86.48</u>	97.55
+H	84.52	79.55	76.15	<u>83.47</u>	85.82	97.60
+A	84.28	78.95	75.97	83.17	85.78	97.55

Multiple-group prediction and suppression. Table III extends this analysis to multi-group attribute prediction and selective suppression. Combining multiple facial-region groups generally improves performance compared to single-group settings, with the strongest improvements obtained when jointly predicting PERIOCLAR, MOUTH, and NOSE. These regions correspond to different facial areas and provide complementary supervisory information, enabling the model to capture a broader range of facial information that is relevant to identity. However, predicting all attribute groups does not further improve performance. In particular, adding HAIR and ACCESSORIES to the {PERIOCLAR, MOUTH, and NOSE} configuration degrades average accuracy.

TABLE III
VERIFICATION ACCURACY (%) WITH MULTIPLE-GROUP ATTRIBUTE SUPERVISION. FR DENOTES THE FACE RECOGNITION BASELINE TRAINED WITHOUT ATTRIBUTE SUPERVISION. P/M/N/H/A DENOTE PERIOCLAR/MOUTH/NOSE/HAIR/ACCESSORIES GROUPS. “+” DENOTES PREDICTION OF THE LISTED ATTRIBUTE GROUPS, WHILE “-” DENOTES SUPPRESSION OF THE SPECIFIED GROUP. AVG. CORRESPONDS TO THE AVERAGE VERIFICATION ACCURACY ACROSS BENCHMARKS.

Mode	Avg.	ECAF _{c2c}	ECAF _{a2c}	AgeDB	CALFW	CACD-VS
FR	84.13	79.35	75.08	81.97	86.58	97.65
+PN	84.31	79.75	75.90	82.98	85.82	97.08
+PM	84.59	79.95	76.08	83.40	86.08	97.42
+PMN	<u>84.66</u>	80.50	<u>76.17</u>	83.08	86.10	<u>97.47</u>
+PMNHA	84.26	79.50	75.35	<u>83.20</u>	85.85	97.42
+PMN-A	84.74	<u>80.30</u>	76.53	83.12	86.08	97.65
+PMN-H	84.04	78.35	75.75	82.83	<u>86.30</u>	96.95

Selective suppression provides further insight: suppressing ACCESSORIES improves performance, supporting the interpretation that accessory-related supervision is weakly tied to identity and can introduce noise. In contrast, suppressing HAIR leads to a large performance drop, indicating that the model relies on hair-related information during training, consistent with the single-group setting where HAIR improved performance despite its higher variability. This reliance likely

reflects dataset-specific attribute–identity correlations, which may inflate performance on the evaluation benchmarks and raise concerns about shortcut learning. We also hypothesize that the smaller improvements on CALFW and CACD-VS reflect their more limited age variation, and that benchmarks with larger age gaps, such as AgeDB and ECAF, benefit more from attribute-based supervision.

Overall, these results establish that attribute supervision can improve face recognition, but that its effectiveness depends strongly on how attributes are grouped and selected. Some groups provide complementary identity information, while others encode dataset-specific correlations that the model may learn to rely on. Selective suppression helps reduce this reliance and enables more controlled use of attribute supervision. These findings highlight the need to analyze attribute groups separately and in combination, rather than treating attribute supervision as a single uniform auxiliary signal.

B. Analysis of Identity Predictiveness in Attribute Labels

The group-wise experiments suggest that attribute supervision can influence recognition not only through complementary identity information, but also through dataset-dependent correlations that the model may exploit. In particular, the strong sensitivity to HAIR suppression motivates a closer examination of how predictive the attribute labels themselves are of identity. To isolate the signal carried by the attribute annotations, we train classifiers using only attribute labels (no images) under two settings: identification and verification. In both settings, we use the MAAD-Face [17] training split and the corresponding attribute annotations.

Identification: We perform closed-set identity classification from attribute vectors using a linear multiclass SVM and report Rank-1 and Rank-5 accuracy.

Verification: We construct an attribute-only verification protocol with an equal number of positive/negative pairs, train a linear binary SVM on 1.2M pairs, and evaluate on a disjoint set of 300k pairs.

Table IV reports identification and verification performance for classifiers trained only on attribute labels, under different single-group and multi-group configurations. Performance is substantially above chance across all groups, reaching 85.31% verification accuracy when all attribute groups are combined, which shows that the attribute annotations alone contain strong identity-related information. Among individual groups, HAIR achieves the highest verification accuracy (79.81%) and the strongest identification results, indicating a particularly strong association between hair-related attributes and identity. These results help interpret the group-wise supervision findings by showing that some attribute groups already carry substantial identity information even without any visual input. In particular, HAIR is highly predictive in the attribute-only setting, which aligns with its strong impact in the recognition experiments: including HAIR can boost performance, and suppressing it causes a marked drop.

TABLE IV

IDENTIFICATION AND VERIFICATION RESULTS USING ATTRIBUTE LABELS ONLY. RANK-1/RANK-5 IDENTIFICATION AND VERIFICATION ACCURACY ARE REPORTED FOR CLASSIFIERS TRAINED WITHOUT IMAGES AND EVALUATED ACROSS ATTRIBUTE GROUPS AND THEIR COMBINATIONS. CHANCE PERFORMANCE IS SHOWN FOR REFERENCE.

Group	Verif. (%)	Rank-1 (%)	Rank-5 (%)
Hair	79.81	<u>1.65</u>	<u>6.12</u>
Periocular	77.32	0.91	3.54
Mouth	75.26	0.53	2.20
Nose	75.55	0.43	1.75
Accessories	75.07	0.45	1.83
P,N,M	<u>81.47</u>	1.56	5.68
All Groups	85.31	4.65	14.05
Chance	50.00	0.02	0.10

Overall, results suggest that attribute supervision effectively improves face recognition performance, while also revealing that some gains may be dataset-dependent and influenced by attribute–identity correlations. The attribute-only analysis confirms that the annotations themselves can be strongly identity-predictive, highlighting the importance of careful attribute selection and control when using attributes for supervision.

VI. CONCLUSION

In this paper, we introduced an attribute-aware face recognition architecture with controlled attribute supervision during training. Attributes are organized into interpretable region-based groups that can be selectively predicted or suppressed, enabling control over how attribute information influences identity learning. Our experiments show that: 1) selective attribute supervision can consistently improve face recognition performance; 2) combining a small complementary set of attributes is more effective and consistent than supervising all groups jointly; and 3) selective suppression via gradient reversal can also reveal contrasting group contributions, as suppressing accessories-related attributes improves performance, while suppressing hair-related attributes causes substantial drops. Our results also show that annotations themselves carry strong identity-related information even without visual input, highlighting that some gains from attribute supervision may be influenced by dataset-dependent attribute–identity correlations. The proposed approach guides models toward identity-relevant information and can reduce reliance on non-discriminative or dataset-specific signals, improving recognition performance and interpretability.

ACKNOWLEDGMENTS

This work was funded by the Portuguese Foundation for Science and Technology (FCT) under the PhD grant 2024.02574.BDANA. This work is supported by NOVA LINCS (UID/04516/2025) with the financial support of FCT.IP (DOI: <https://doi.org/10.54499/UID/04516/2025>). It was also supported, when eligible, by national funds through FCT under project UID/50008/2025 – Instituto de Telecomunicações (DOI: <https://doi.org/10.54499/UID/50008/2025>).

REFERENCES

- [1] H. Wang, Y. Wang *et al.*, “CosFace: Large margin cosine loss for deep face recognition,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2018.
- [2] J. Deng, J. Guo *et al.*, “ArcFace: Additive angular margin loss for deep face recognition,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 10, pp. 5962–5979, 2022.
- [3] Q. Meng, S. Zhao *et al.*, “MagFace: A Universal Representation for Face Recognition and Quality Assessment,” *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 14 220–14 229, 2021.
- [4] Z. Huang, J. Zhang *et al.*, “When age-invariant face recognition meets face age synthesis: A multi-task learning framework and a new benchmark,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 6, pp. 7917–7932, 2023.
- [5] J.-C. Xie, C.-M. Pun *et al.*, “Implicit and explicit feature purification for age-invariant facial representation learning,” *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 399–412, 2022.
- [6] K. Cao, Y. Rong *et al.*, “Pose-robust face recognition via deep residual equivariant mapping,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2018.
- [7] Y. Liu, W. Wang *et al.*, “Pose-disentangled contrastive learning for self-supervised facial representation,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2023, pp. 9717–9728.
- [8] L. Song, D. Gong *et al.*, “Occlusion robust face recognition based on mask learning with pairwise differential siamese network,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2019, pp. 773–782.
- [9] T. Feng, L. Xu *et al.*, “Towards mask-robust face recognition,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2021, pp. 1492–1496.
- [10] F. Taherkhani, N. M. Nasrabadi *et al.*, “A deep face identification network enhanced by facial attributes prediction,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, 2018.
- [11] Z. Wang, K. He *et al.*, “Multi-task deep neural network for joint face recognition and facial attribute prediction,” in *Proc. ACM Int. Conf. Multimedia Retrieval (ICMR)*, 2017, pp. 365–374.
- [12] L. Qin, M. Wang *et al.*, “SwinFace: A multi-task transformer for face recognition, expression recognition, age estimation and attribute estimation,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 34, no. 4, p. 2223–2234, 2024.
- [13] Qin, Lixiong, M. Wang *et al.*, “Facepor: A generalist model for face perception,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Springer, 2024, pp. 240–260.
- [14] Y. Zhong, J. Sullivan *et al.*, “Face attribute prediction using off-the-shelf cnn features,” in *Proc. Int. Conf. Biometrics (ICB)*, 2016, pp. 1–7.
- [15] Y. Ganin, E. Ustinova *et al.*, “Domain-adversarial training of neural networks,” *J. Mach. Learn. Research*, vol. 17, no. 59, pp. 1–35, 2016.
- [16] H. Wang, V. Sanchez *et al.*, “From age estimation to age-invariant face recognition: Generalized age feature extraction using order-enhanced contrastive learning,” *IEEE Trans. Inf. Forensics Security*, 2025.
- [17] P. Terhörst, D. Fährmann *et al.*, “Maad-face: A massively annotated attribute dataset for face images,” *IEEE Trans. Inf. Forensics Security*, vol. 16, pp. 3942–3957, 2021.
- [18] Z. Liu, P. Luo *et al.*, “Deep learning face attributes in the wild,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2015.
- [19] J. Deng, J. Guo *et al.*, “RetinaFace: Single-shot multi-level face localization in the wild,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2020.
- [20] K. Zhang, Z. Zhang *et al.*, “Joint face detection and alignment using multitask cascaded convolutional networks,” *IEEE Signal Process. Lett.*, vol. 23, no. 10, pp. 1499–1503, 2016.
- [21] T. Zheng, W. Deng *et al.*, “Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments,” *arXiv preprint arXiv:1708.08197*, 2017.
- [22] S. Moschoglou, A. Papaioannou *et al.*, “AgeDB: The first manually collected, in-the-wild age database,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, 2017, pp. 1997–2005.
- [23] B.-C. Chen, C.-S. Chen *et al.*, “Cross-age reference coding for age-invariant face recognition and retrieval,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2014, pp. 768–783.
- [24] Z. Liu, Y. Lin *et al.*, “Swin transformer: Hierarchical vision transformer using shifted windows,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2021, pp. 10 012–10 022.