

ArtFace: Towards Historical Portrait Face Identification via Model Adaptation

Francois Poh, Anjith George and Sébastien Marcel
Idiap Research Institute
Rue Marconi 19, CH - 1920, Martigny, Switzerland

francois.poh22@imperial.ac.uk, {anjith.george, marcel}@idiap.ch

1. Introduction

Identifying sitters in historical paintings is a key task for art historians, offering insight into their lives and how they chose to be seen [11]. However, the process is often subjective and limited by the lack of data and stylistic variations. Automated facial recognition is capable of handling challenging conditions and can assist, but while traditional facial recognition models perform well on photographs [7], they struggle with paintings due to domain shift and high intra-class variation [5]. Artistic factors such as style, skill, intent [8, 11], and influence from other works [1] further complicate recognition.

The early work of Srinivasan et al. [14] approached the problem using local and anthropometric facial features to statistically assess similarity between portrait pairs. More recent approaches rely heavily on deep learning, particularly convolutional neural networks (CNNs). Gupta et al. [5] used Siamese CNNs to compare controversial historical portraits, leveraging style transfer techniques to augment training data while preserving essential facial content. Building on this, Huber et al. [8] fine-tuned a ResNet-100 model on an augmented dataset and introduced a confidence measure by estimating uncertainty in face comparison scores. Salavessa et al. [11] combined deep learning with handcrafted facial feature ratios, using a VGG-based approach.

In this work, we investigate the potential of foundation models to improve facial recognition in artworks. By fine-tuning foundation models and integrating their embeddings with those from conventional facial recognition networks, we demonstrate notable improvements over current state-of-the-art methods. Our results show that foundation models can bridge the gap where traditional methods are ineffective. Paper page at ¹.

Contributions We address the task of sitter identification through the following contributions: (i) Fine-tuning the vision-language foundation model CLIP [10] with LoRA [6]

on portrait face images; (ii) Complementing it with embeddings from an adapted face recognition network (AntelopeV2/IResNet100 [2]); and (iii) Fusing the embeddings via a simple yet effective normalised concatenation strategy. Comprehensive experiments on the public *Historical Faces* dataset [8] show that our method reduces baseline error rates and outperforms prior approaches.

2. Proposed Method

CLIP Foundation Model: CLIP is a vision-language foundation model trained on 400 million image-text pairs collected from the internet [10]. It learns to align images and text in a shared embedding space, enabling zero-shot transfer to a wide range of tasks such as image classification and visual understanding. Since foundation models are trained on diverse data, they capture broad concepts, detect stylistic differences, and can use contextual information relevant to paintings that traditional facial recognition overlooks. Prior work has shown that combining CLIP embeddings with those from facial recognition models yields modest improvements on photographic data [13], and this benefit can extend to painted portraits as well. For our experiments, we use the CLIP ViT-B/16 model from OpenAI, selected for its balance between tunability and computational efficiency.

Face Recognition Model: AntelopeV2 is a state-of-the-art face recognition (FR) model built on the IResNet100 architecture [2] and trained on the large-scale Glint360k dataset. It achieves high accuracy and robustness across challenging variations in pose, age, and lighting, making it a strong and widely adopted baseline in face verification benchmarks. In addition to these models, we also use a Commercial off-the-shelf face recognition model as another baseline (COTS).

Model Adaptation: Both the CLIP and face recognition (FR) models are adapted using portrait images to better align with the domain of historical paintings, as shown in Fig. 1. For CLIP, we apply Low-Rank Adaptation (LoRA) [6], inserting LoRA layers into the query (Q) and value (V) matrices of the attention mechanism. This ap-

¹<https://www.idiap.ch/paper/artface/>

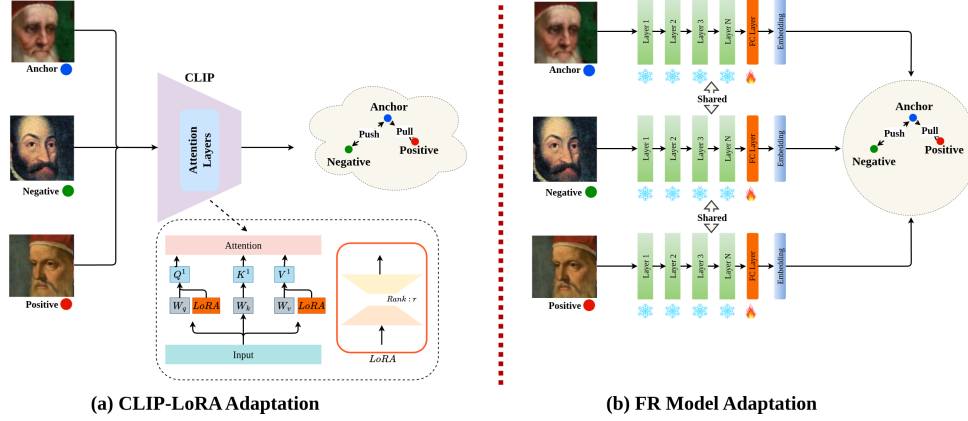


Figure 1. Overview of the proposed method: (a) LoRA-based adaptation of the CLIP model, and (b) head adaptation using triplet loss.

proach introduces low-rank updates, allowing efficient fine-tuning with significantly fewer trainable parameters and reduced memory usage, while maintaining performance. For the IResNet100 model, adaptation is minimal; only the final linear layer is fine-tuned to adjust the embeddings for the new domain.

Fusion: Embeddings from the tuned and untuned IResNet100 models, along with the tuned CLIP model, are first individually normalised, then concatenated and re-normalised. This process effectively performs feature fusion, combining complementary information from each model.

3. Experiments

Dataset: The Historical Faces dataset [8] consists of 766 paintings of 210 different sitters. The paintings are pre-processed by detecting facial key points and aligning and cropping to 112x112 for both IResNet100 and CLIP. This was split for training and evaluation by identity in a 60:20:20 ratio for the train, validation, and test set, respectively.

Training Details: The CLIP-LoRA model is trained using triplet loss [12] with cosine distance and a margin of 0.5, with a batch size of 48. Training employs early stopping with a patience of 10 epochs to prevent overfitting. Hard negative mining is performed at each epoch: 30% of the negative samples are selected from the top 50 hardest negatives per anchor, while the remaining 70% are randomly drawn from the next 450 hardest examples. The model is optimised using Adam with a learning rate of 1e-5. LoRA layers, with rank 16, are applied to the query (Q) and value (V) matrices in all transformer layers, implemented using the PEFT library [9]. IResNet100 uses the same method and parameters for training, but no LoRA layers are applied; instead, only tuning the final linear layer. The final output score is the cosine similarity between the concatenated embeddings.

3.1. Experimental Results

Results with IResNet100 Tuning: Table 1 demonstrates that tuning IResNet100 primarily enhances TAR² at low FAR³ values (<1%), with only modest improvements in EER⁴. While fusion performs slightly worse than the tuned model in low-FAR TAR, it consistently outperforms both the base and tuned models overall.

Table 1. Performance of IResNet100 variants with and without tuning and fusion.

Model	EER↓	TAR@0.1%FAR↑	TAR@1%FAR↑
IResNet100-Base	14.0%	29.9%	55.1%
IResNet100-Tuned	13.5%	36.5%	53.7%
IResNet100-Fusion (Base+Tuned)	13.5%	33.6%	58.4%

Results with CLIP-LoRA Tuning: In this section, we describe the fine-tuning of CLIP-LoRA using triplet loss, both with and without hard negative mining. All other relevant parameters remain unchanged from previous experiments. Although the hard negative mining variant does not achieve higher TAR at all FAR levels, it shows improved performance above approximately 1% FAR. As shown in Table 2, this approach also leads to better EER.

Results with Fusion: Table 3 compares the performance of fusion methods with individual IResNet100 and CLIP models on the Historical Faces test split. Although the

²True Accept Rate (TAR) = 1 - False Non-Match Rate (FNMR)

³False Accept Rate (FAR) = False Match Rate (FMR)

⁴Equal Error Rate (EER) = Error Rate when FMR=FNMR

Table 2. Performance of loss functions for tuning CLIP. Models with Hard Negative Mining (HN) perform better overall.

Model	EER↓	TAR@0.1%FAR↑	TAR@1%FAR↑
CLIP-Base	17.9%	8.4%	33.2%
CLIP-LoRA(Triplet, HN)	13.1%	17.8%	43.5%
CLIP-LoRA(Triplet, w/o HN)	13.9%	16.8%	43.9%

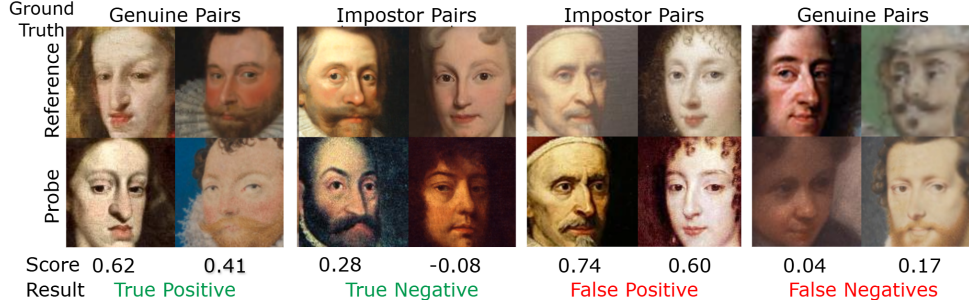


Figure 2. Examples of successful and failed comparisons. Each pair shows the reference and probe images, with the associated cosine similarity score. High similarity scores indicate successful matches, while low scores represent mismatches.

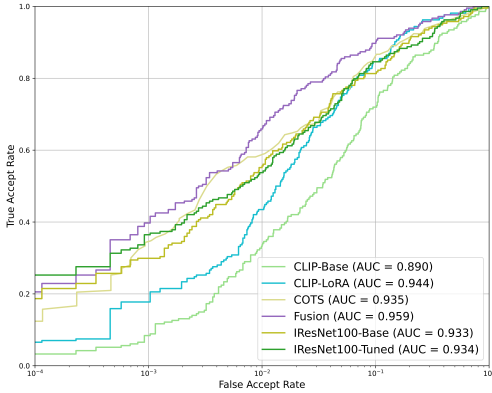


Figure 3. ROC curves of tuned and base CLIP, IResNet100, COTS and proposed fusion method. Fusion provides consistent improvements even at low FAR.

Table 3. Performance Comparison of Base, Tuned models, Fusion, and COTS FR Systems. Fusion enhances overall accuracy.

Model	EER↓	TAR@0.1%FAR↑	TAR@1%FAR↑
COTS FR system	12.6%	34.3%	58.1%
CLIP-Base	17.9%	8.4%	33.2%
IResNet100-Base	14.0%	29.9%	55.1%
CLIP-Base + IResNet100-Base	13.1%	29.0%	54.7%
CLIP-Base + IResNet100-Tuned	12.6%	35.1%	57.9%
CLIP-LoRA + IResNet100-Base	11.1%	34.6%	62.6%
CLIP-LoRA + IResNet100-Tuned	10.7%	39.7%	62.15%
CLIP-LoRA + IResNet100-Base + IResNet100-Tuned	9.9%	39.7%	65.9%

base CLIP model underperforms compared to IResNet100, fusion leads to improvements in EER. Fine-tuning CLIP-LoRA further enhances performance across all four metrics, with additional gains observed when the tuned IResNet100 is included in the fusion. The corresponding ROC curves are presented in Fig. 3.

Discussions: Tuning IResNet100 on the Historical Faces dataset yields only modest gains, but fusing tuned and untuned versions outperforms either alone, suggesting combining representations can help, even within a single architecture.

CLIP-LoRA tuning shows triplet loss with hard negative

mining performs best, significantly improving over base CLIP. Other loss functions underperform, likely due to limited data, especially in terms of images per identity.

Fusion experiments show that combining foundation models with conventional face recognition networks improves performance across all evaluation metrics, especially in low false acceptance rate (FAR) scenarios. This indicates that foundation models capture valuable information from portraits that conventional networks may overlook.

Visualisations: To better understand both successful and failed matches, we show sample images along with their similarity scores in Fig. 2. Genuine comparisons can fail in challenging cases, particularly when there are significant differences in artistic style or medium, sitter age, and often compounded when multiple factors overlap. Conversely, some impostor comparisons yield surprisingly high similarity scores. This typically occurs when artists base their work on earlier portraits or when sitters resemble each other and are depicted in a similar style or by the same artist.

4. Conclusions

In this work, we show that lightweight tuning of vision-language foundation models [6], combined with domain-adapted face recognition networks, can effectively bridge the domain gap between photographs and paintings. Our fusion approach achieves state-of-the-art accuracy in sitter identification. Face recognition on artworks remains a particularly difficult task compared to traditional FR due to the scarcity of labelled data, stylistic variation, and the interpretive nature of portraiture. However, the results show that adapting modern architectures to this setting is feasible and promising. This opens up new research avenues, including synthetic data generation [4] to augment the limited training set and heterogeneous domain adaptation techniques [3] to improve generalisation across visual domains.

5. Acknowledgements

This research was supported by the InterArt project, funded by Loterie Romande.

References

- [1] Art UK. Why artists copy: A history of drawing the masters, 2024. Accessed: 2025-07-25.
- [2] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019.
- [3] Anjith George and Sébastien Marcel. Modality agnostic heterogeneous face recognition with switch style modulators. In *2024 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10. IEEE, 2024.
- [4] Anjith George and Sébastien Marcel. Digi2real: Bridging the realism gap in synthetic data face recognition via foundation models. In *Proceedings of the Winter Conference on Applications of Computer Vision*, pages 1469–1478, 2025.
- [5] Akash Gupta, Niluthpol C Mithun, Conrad Rudolph, and Amit K Roy-Chowdhury. Deep learning based identity verification in renaissance portraits. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2018.
- [6] 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*, 1(2):3, 2022.
- [7] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*, 2008.
- [8] Marco Huber, Philipp Terhörst, Anh Thi Luu, Florian Kirchbuchner, and Naser Damer. Verification of sitter identity across historical portrait paintings by confidence-aware face recognition. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pages 938–944. IEEE, 2022.
- [9] Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin Bossan. PEFT: State-of-the-art parameter-efficient fine-tuning methods. <https://github.com/huggingface/peft>, 2022.
- [10] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021.
- [11] Eunice Salavessa, José Aranha, Rafael Moreira, and David M Freire-Lista. Proto-early renaissance depictions, iconographic analysis and computerised facial similarity assessment connections: The 16th century mural paintings of st. leocadia church (chaves, north of portugal). *Heritage*, 7(4):2031–2054, 2024.
- [12] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [13] Redwan Sony, Parisa Farmanifard, Arun Ross, and Anil K Jain. Foundation versus domain-specific models: Performance comparison, fusion, and explainability in face recognition. *arXiv preprint arXiv:2507.03541*, 2025.
- [14] Ramya Srinivasan, Conrad Rudolph, and Amit K. Roy-Chowdhury. Computerized face recognition in renaissance portrait art: A quantitative measure for identifying uncertain subjects in ancient portraits. *IEEE Signal Processing Magazine*, 32(4):85–94, 2015.