

# Beyond Spatial Explanations: Explainable Face Recognition in the Frequency Domain - Supplementary Material

Marco Huber and Naser Damer  
Fraunhofer Institute for Computer Graphics Research IGD, Germany  
Technische Universität Darmstadt, Germany  
Email: marco.huber@igd.fraunhofer.de

## 1. Introduction

This is the supplementary material to the paper: *Beyond Spatial Explanations: Explainable Face Recognition in the Frequency Domain*. In this supplementary material, we provide additional quantitative results, the approach applied to the use case of morphing attacks, and more visual examples and investigations. While we provided the insertion and deletion curves using the  $\mathcal{L}_2$ -norm following Wang *et al.* [48] as the norm in Equation (2) in the paper for the LFW dataset [14], we provide here the insertion and deletion curves using the  $\mathcal{L}_1$ -norm as proposed by Abello *et al.* [1]. Furthermore, we provide the quantitative results on the Cross-Age LFW dataset [51], which provides a more challenging scenario as the LFW dataset [14]. For the visual investigation and examples, we provide additional samples using different bandsizes, the two different face recognition models (ElasticFace-Arc [4], CurricularFace [15]), and the two different norms ( $\mathcal{L}_1$ ,  $\mathcal{L}_2$ ).

## 2. Qualitative Results using $\mathcal{L}_2$ -norm

Due to space we only provided the qualitative results masking utilizing  $\mathcal{L}_2$ -norm as proposed by Wang *et al.* [48]. In contrast, Abello *et al.* [1] proposed to use  $\mathcal{L}_1$ -norm rather than  $\mathcal{L}_2$ -norm for masking in the frequency domain due to the better suitability for discrete spaces [1]. The results are provided in Fig. 1 for ElasticFace-Arc [4] and in Fig. 2 for CurricularFace [15]. Similar to the results presented in the paper using  $\mathcal{L}_2$ -norm, the faster increasing deletion curves and the faster increasing insertion curves prove the effectiveness of our approach also when using  $\mathcal{L}_1$ -norm with only a slightly difference between the different norms.

### 2.1. Use Case: Morphing Attacks

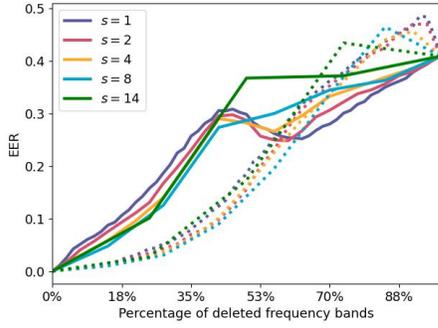
As a second use case in addition to the provided cross-resolution use case presented in the main paper, we investigate explaining the frequency importance of bona fide and morphing attack comparison to references with the assumption that the morphing process would effect some fre-

quency component of the image, and thus its comparison to the reference. For the morphing attack images, we utilize the SYN-MAD 2022 dataset [55] which consists of around 1000 morphing images of each morphing approach and 204 bona fide images. In Figure 7, we provide the mean absolute FHPs plot and the standard deviation for the FHPs with  $s = 8$ . The figures shows that the frequency importance is wider distributed for all three morphing attack types (OpenCV (landmark-based), MIPGAN2 (GAN-based) [56], MorDIFF (diffusion-based) [54]) in comparison to the bona fide FHPs and also show higher variety (shown as the error bar, standard deviation). This allows us to conclude based on the proposed frequency-based explanations, that FR models process morphing attack differently than unaltered bona fide images. These findings support the usability of the frequency-based explanations in diverse use cases to gain new insights into FR model behavior.

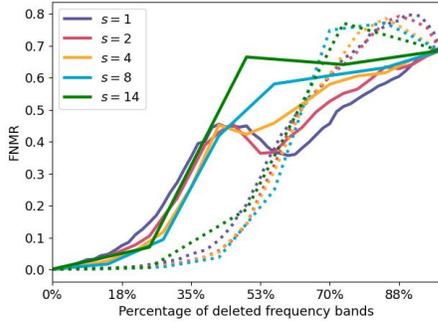
### 2.2. Quantitative Results on Cross-Age LFW

As a second and more challenging dataset, we selected Cross-Age LFW (CALFW) [43]. The CALFW dataset has been collected with the aim to better represent age gaps in the genuine pairs, leading to a more challenging face recognition scenario. The preprocessing follows the same procedure as the preprocessing reported in the main paper on the LFW dataset. We report the quantitative results in terms of insertion and deletion curves for both,  $\mathcal{L}_1$ -norm and  $\mathcal{L}_2$ -norm, and also for both, EER and FNMR as the evaluation point.

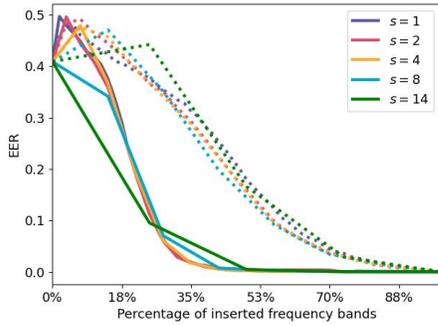
The results using ElasticFace-Ar [4] are provided in Figure 3 ( $\mathcal{L}_1$ -norm) and Figure 4 ( $\mathcal{L}_2$ -norm). The results using CurricularFace [15] are provided in Figure 5 ( $\mathcal{L}_1$ -norm) and Figure 6 ( $\mathcal{L}_2$ -norm). The results on Cross-age LFW show, that the proposed approach also generalizes to more challenging datasets as a similar performance can be observed compared to the results obtained on LFW (Figure 3 and 4 in the main paper ( $\mathcal{L}_2$ -norm) and Figure 1 and 2 ( $\mathcal{L}_1$ -norm).



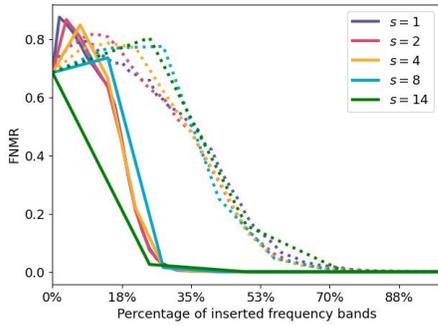
(a) Deletion - EER



(b) Deletion - FNMR

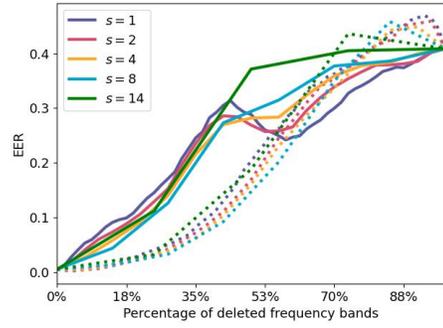


(c) Insertion - EER

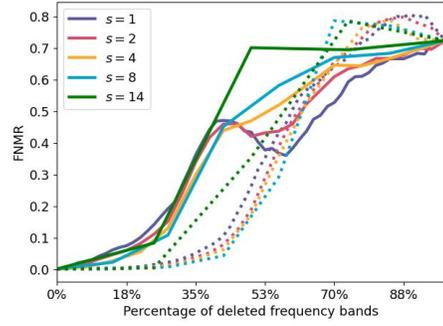


(d) Insertion - FNMR

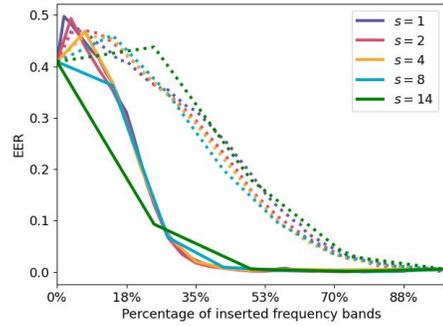
Figure 1. Deletion and insertion curves using ElasticFace-Arc [4] and  $\mathcal{L}_1$ -norm while Figure 3 in the main paper uses  $\mathcal{L}_2$ -norm. The solid lines are a result of our proposed explanations. The dotted line indicates the performance of the baseline with the same frequency band size  $s$  as its color counterpart in solid line. Both faster ascending deletion curves and faster descending insertion curves point the effectiveness of the proposed explanations.



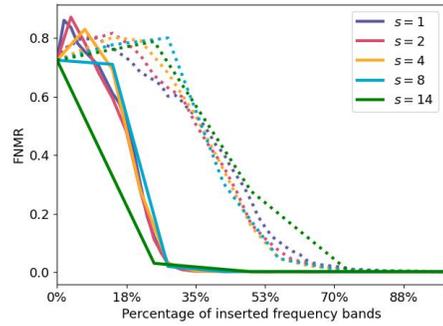
(a) Deletion - EER



(b) Deletion - FNMR

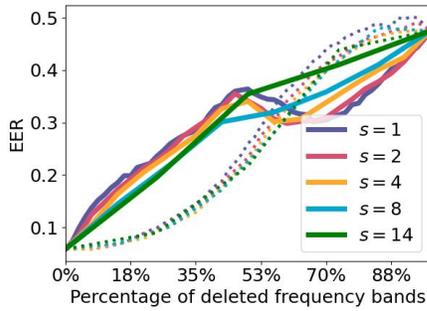


(c) Insertion - EER

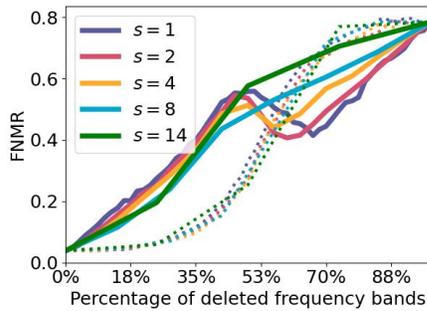


(d) Insertion - FNMR

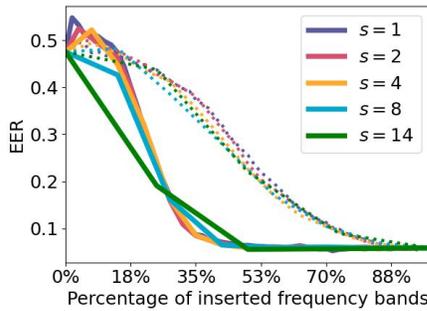
Figure 2. Deletion and insertion curves using CurricularFace [15] and  $\mathcal{L}_1$ -norm while Figure 4 in the main paper uses  $\mathcal{L}_2$ -norm.. The solid lines are a result of our proposed explanations. The dotted line indicates the performance of the baseline with the same frequency band size  $s$  as its color counterpart in solid line. Both faster ascending deletion curves and faster descending insertion curves point the effectiveness of the proposed explanations.



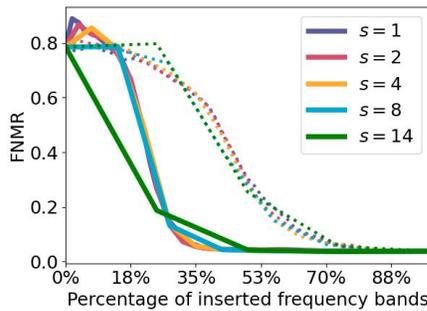
(a) Deletion - EER



(b) Deletion - FNMR

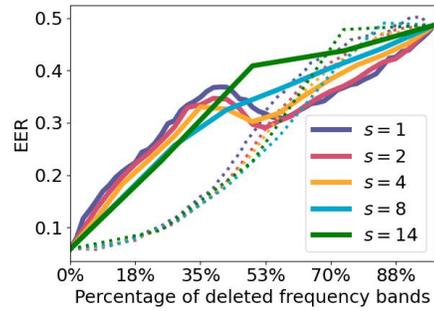


(c) Insertion - EER

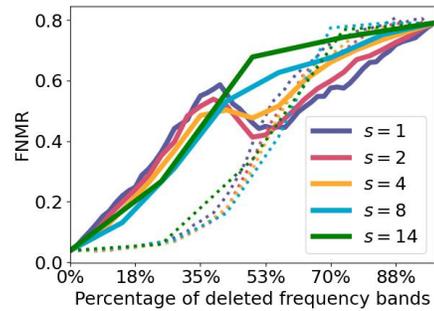


(d) Insertion - FNMR

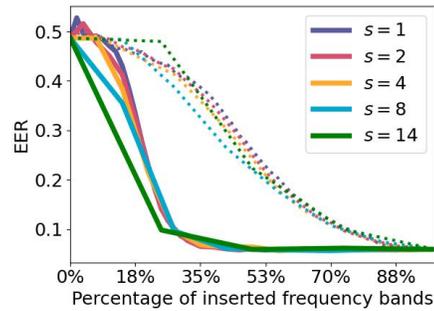
Figure 3. Deletion and insertion curves using ElasticFace-Arc [4] and  $\mathcal{L}_1$ -norm on the Cross-Age LFW [51] dataset. The solid lines are a result of our proposed explanations. The dotted line indicates the performance of the baseline with the same frequency band size  $s$  as its color counterpart in solid line. The obtained results are similar to the results obtained on the LFW [14] dataset, as both faster ascending deletion curves and faster descending insertion curves point the effectiveness of the proposed explanations.



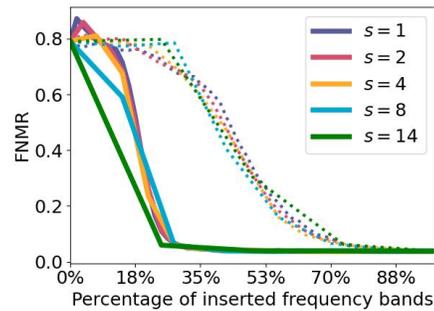
(a) Deletion - EER



(b) Deletion - FNMR

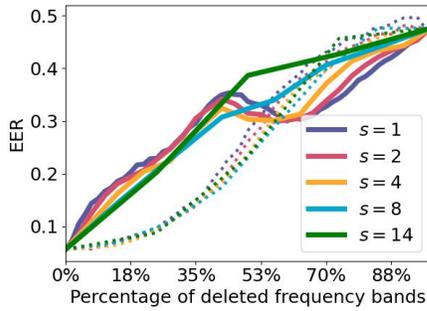


(c) Insertion - EER

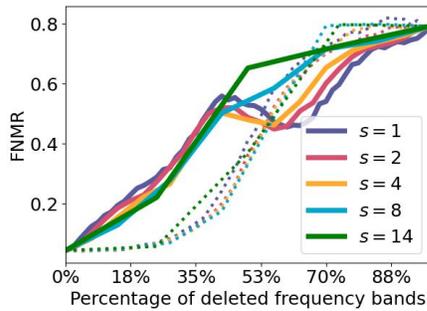


(d) Insertion - FNMR

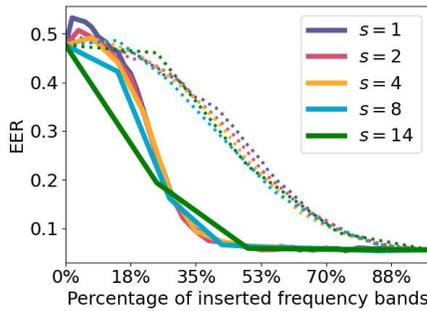
Figure 4. Deletion and insertion curves using ElasticFace-Arc [4] and  $\mathcal{L}_2$ -norm on the Cross-Age LFW [51] dataset. The solid lines are a result of our proposed explanations. The dotted line indicates the performance of the baseline with the same frequency band size  $s$  as its color counterpart in solid line. The obtained results are similar to the results obtained on the LFW [14] dataset, as both faster ascending deletion curves and faster descending insertion curves point the effectiveness of the proposed explanations.



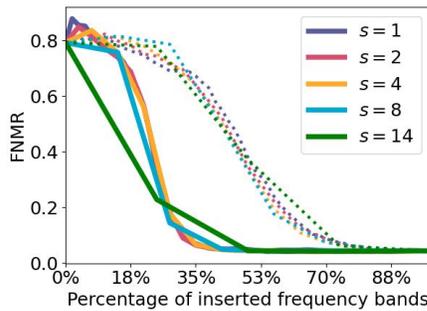
(a) Deletion - EER



(b) Deletion - FNMR

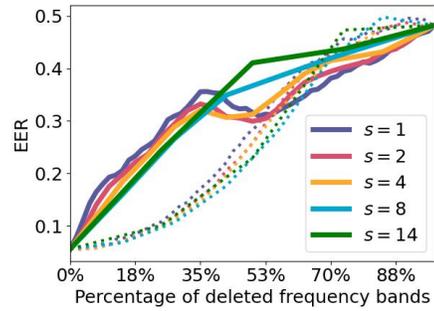


(c) Insertion - EER

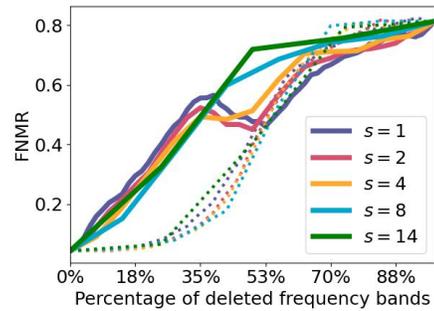


(d) Insertion - FNMR

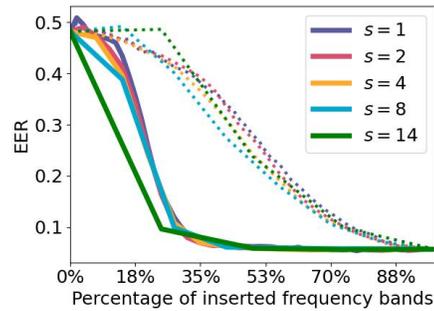
Figure 5. Deletion and insertion curves using CurricularFace [15] and  $\mathcal{L}_1$ -norm on the Cross-Age LFW [51] dataset. The solid lines are a result of our proposed explanations. The dotted line indicates the performance of the baseline with the same frequency band size  $s$  as its color counterpart in solid line. The obtained results are similar to the results obtained on the LFW [14] dataset, as both faster ascending deletion curves and faster descending insertion curves point the effectiveness of the proposed explanations.



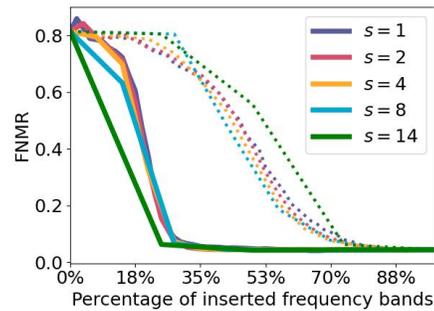
(a) Deletion - EER



(b) Deletion - FNMR



(c) Insertion - EER



(d) Insertion - FNMR

Figure 6. Deletion and insertion curves using CurricularFace [15] and  $\mathcal{L}_2$ -norm on the Cross-Age LFW [51] dataset. The solid lines are a result of our proposed explanations. The dotted line indicates the performance of the baseline with the same frequency band size  $s$  as its color counterpart in solid line. The obtained results are similar to the results obtained on the LFW [14] dataset, as both faster ascending deletion curves and faster descending insertion curves point the effectiveness of the proposed explanations.

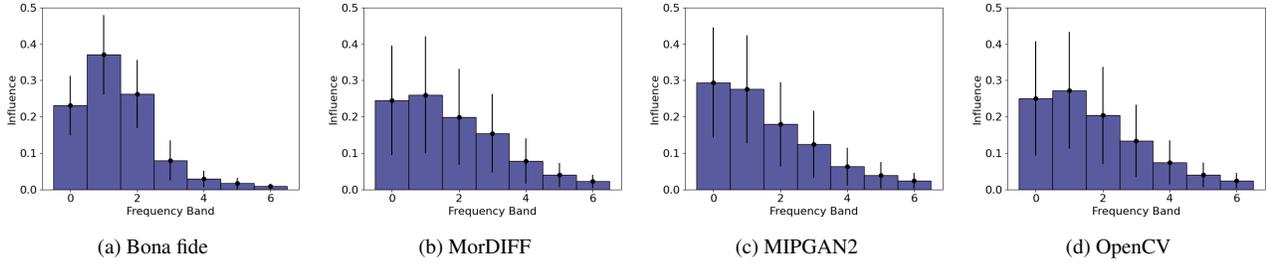


Figure 7. Mean absolute FHPs for bona fide pairs and three different morphing attacks. In morphing attack pairs, the frequency importance is more distributed compared to the bona fide importance and the variety is higher (error bar, standard deviation).

### 3. Visual Investigations

#### 3.1. Comparison of $L_1$ - and $L_2$ -norm

Due to space limitations, we provided in the paper only FHPs using  $\mathcal{L}_2$ -norm during the masking process of our approach. Here we present the absolute and directed Frequency Heat Plots (FHPs) utilizing either  $\mathcal{L}_1$ - or  $\mathcal{L}_2$ -norm on the unaltered and the low-resolution image pairs for both, matches and non-matches. We also include the FHPs for all investigated bandsizes  $s \in [1, 2, 4, 8, 14]$ . The results for the unaltered pairs are presented in Figures 8, 9, 10, and 11. The results on low-resolution image pairs are presented in Figures 12, 13, 14 and 15. Similar to the observation in the quantitative analysis above, the assigned influences for the different frequency bands is quite similar, independent of the utilized norm.

#### 3.2. Comparison between ElasticFace-Arc and CurricularFace

We provide in the paper the quantitative results for both FR models and some FHPs utilizing the ElasticFace-Arc [4] model (due to limited space). Here, we additionally provide some examples as an comparison between the absolute and directed FHPs on matching and non-matching image pairs, as well as unaltered and low-resolution image pairs. The FHPs based on both models for matching pairs are provided in the Figures 16, 17, 18 and 19. For the non-matching pairs, some FHPs based on both models are provided in the Figures 20, 21, 22 and 23. The FHPs provided show, that to some extent, the same frequency bands are similarly influential on the image pairs, independent from the model utilized.

### References

[54] Naser Damer, Meiling Fang, Patrick Siebke, Jan Niklas Kolf, Marco Huber, and Fadi Boutros. Mordiff: Recognition vulnerability and attack detectability of face morphing attacks created by diffusion autoencoders. In *11th International Workshop on Biometrics and Forensics, IWBF 2023, Barcelona, Spain, April 19-20, 2023*, pages 1–6. IEEE, 2023. 1

[55] Marco Huber, Fadi Boutros, Anh Thi Luu, Kiran B. Raja, Raghavendra Ramachandra, Naser Damer, Pedro C. Neto, Tiago Gonçalves, Ana Filipa Sequeira, Jaime S. Cardoso, João Tremoço, Miguel Lourenço, Sergio Serra, Eduardo Cermeño, Marija Ivanovska, Borut Batagelj, Andrej Kronovsek, Peter Peer, and Vitomir Struc. SYN-MAD 2022: Competition on face morphing attack detection based on privacy-aware synthetic training data. In *IEEE International Joint Conference on Biometrics, IJCB 2022, Abu Dhabi, United Arab Emirates, October 10-13, 2022*, pages 1–10. IEEE, 2022. 1

[56] Haoyu Zhang, Sushma Venkatesh, Raghavendra Ramachandra, Kiran Bylappa Raja, Naser Damer, and Christoph Busch. MIPGAN - generating strong and high quality morphing attacks using identity prior driven GAN. *IEEE Trans. Biom. Behav. Identity Sci.*, 3(3):365–383, 2021. 1

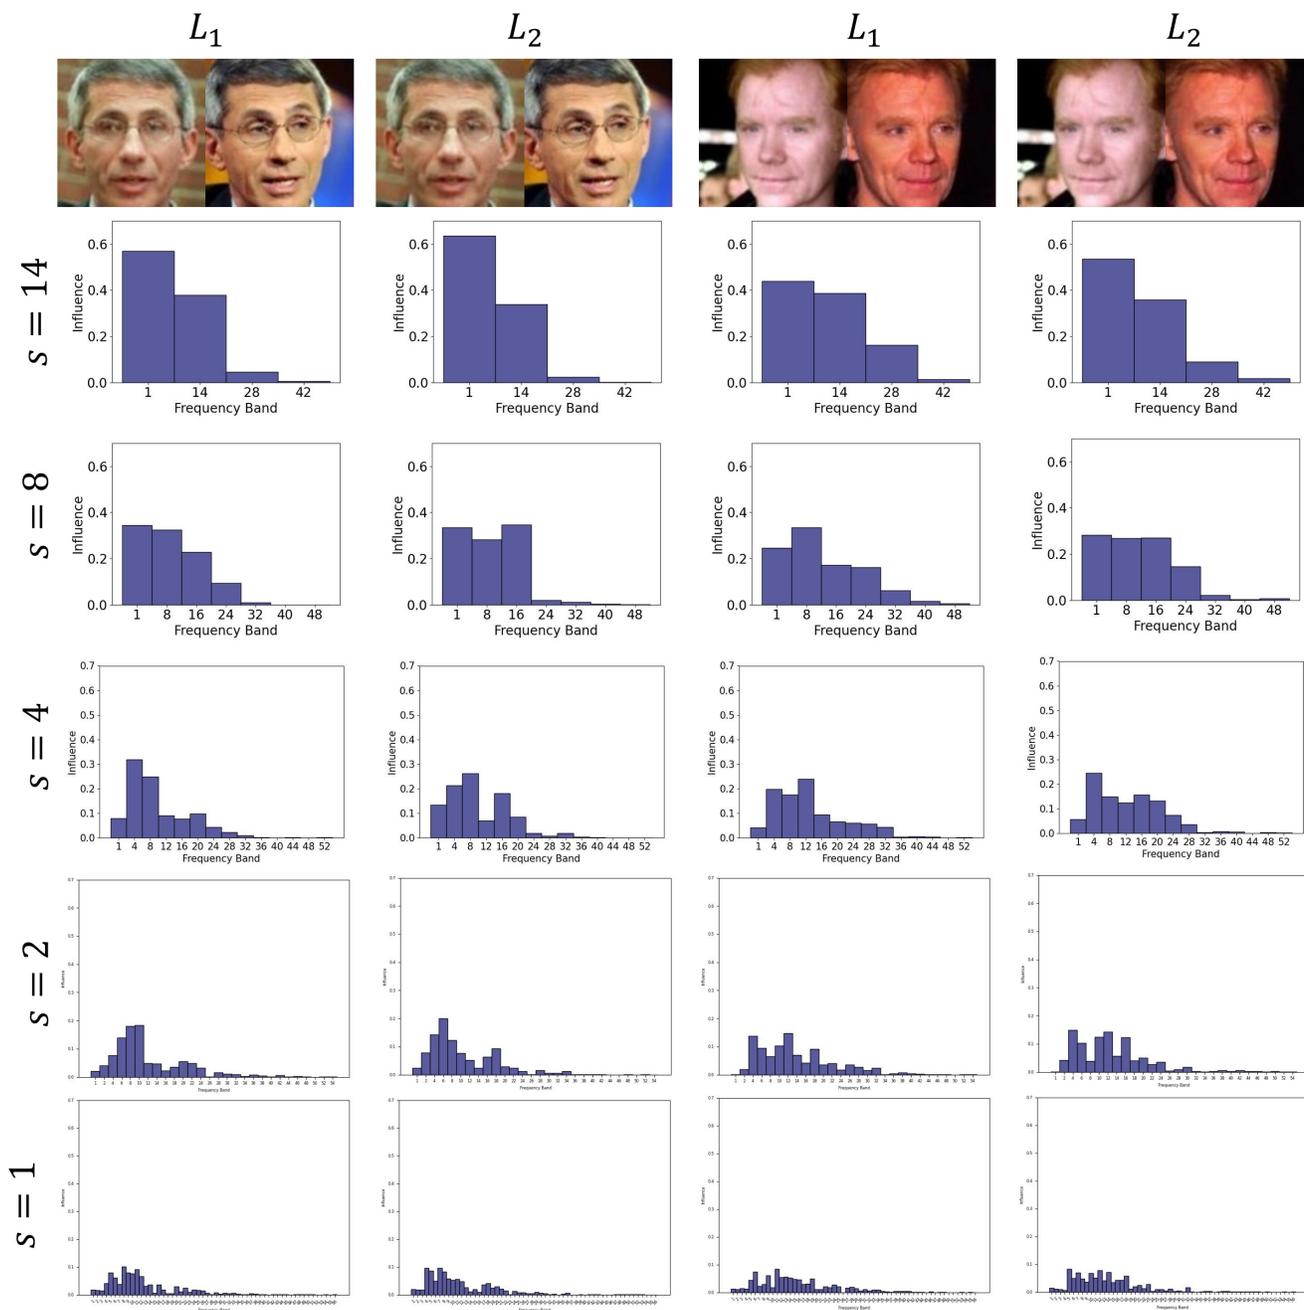


Figure 8. Absolute FHPs using the ElasticFace-Arc [4] model on unaltered matching pairs and either  $\mathcal{L}_1$ - or  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show some similarity independent of the norm utilized for masking, especially with smaller bandsizes. For the sake of comparison, we keep the y-axis scale fixed over all absolute FHPs.

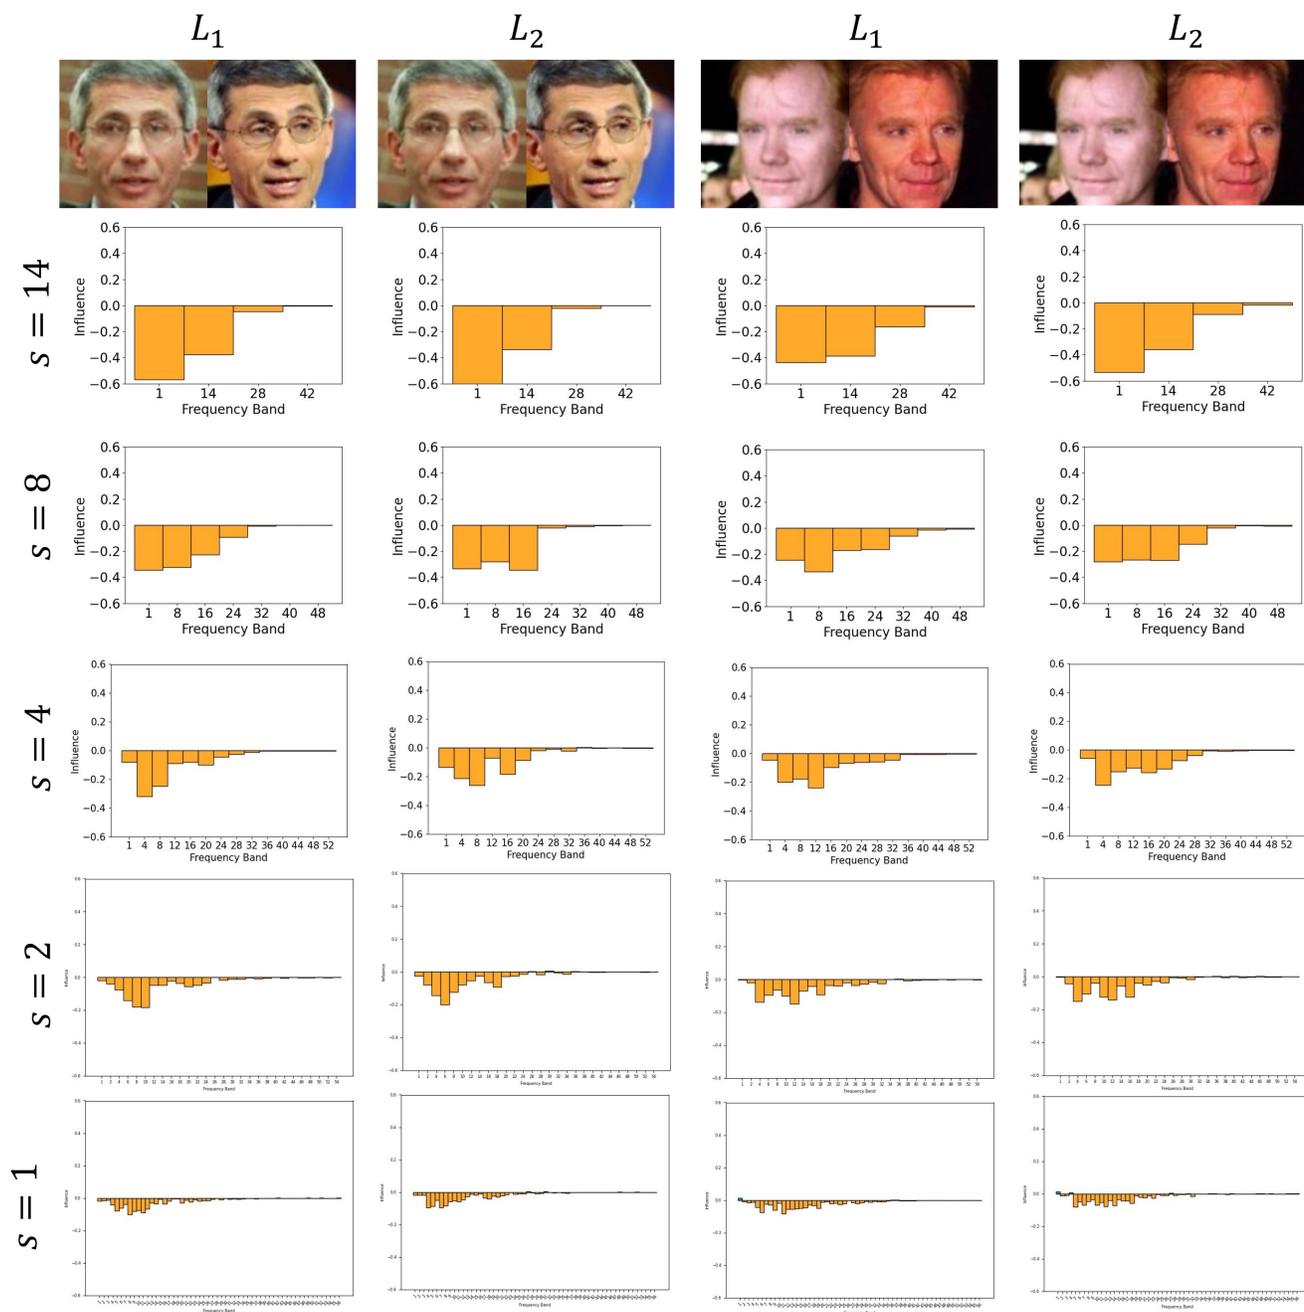


Figure 9. Directed FHPs using the ElasticFace-Arc [4] model on unaltered matching pairs and either  $\mathcal{L}_1$ - or  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show some similarity independent of the norm utilized for masking, especially with smaller bandsizes. For the sake of comparison, we keep the y-axis scale fixed over all directed FHPs.

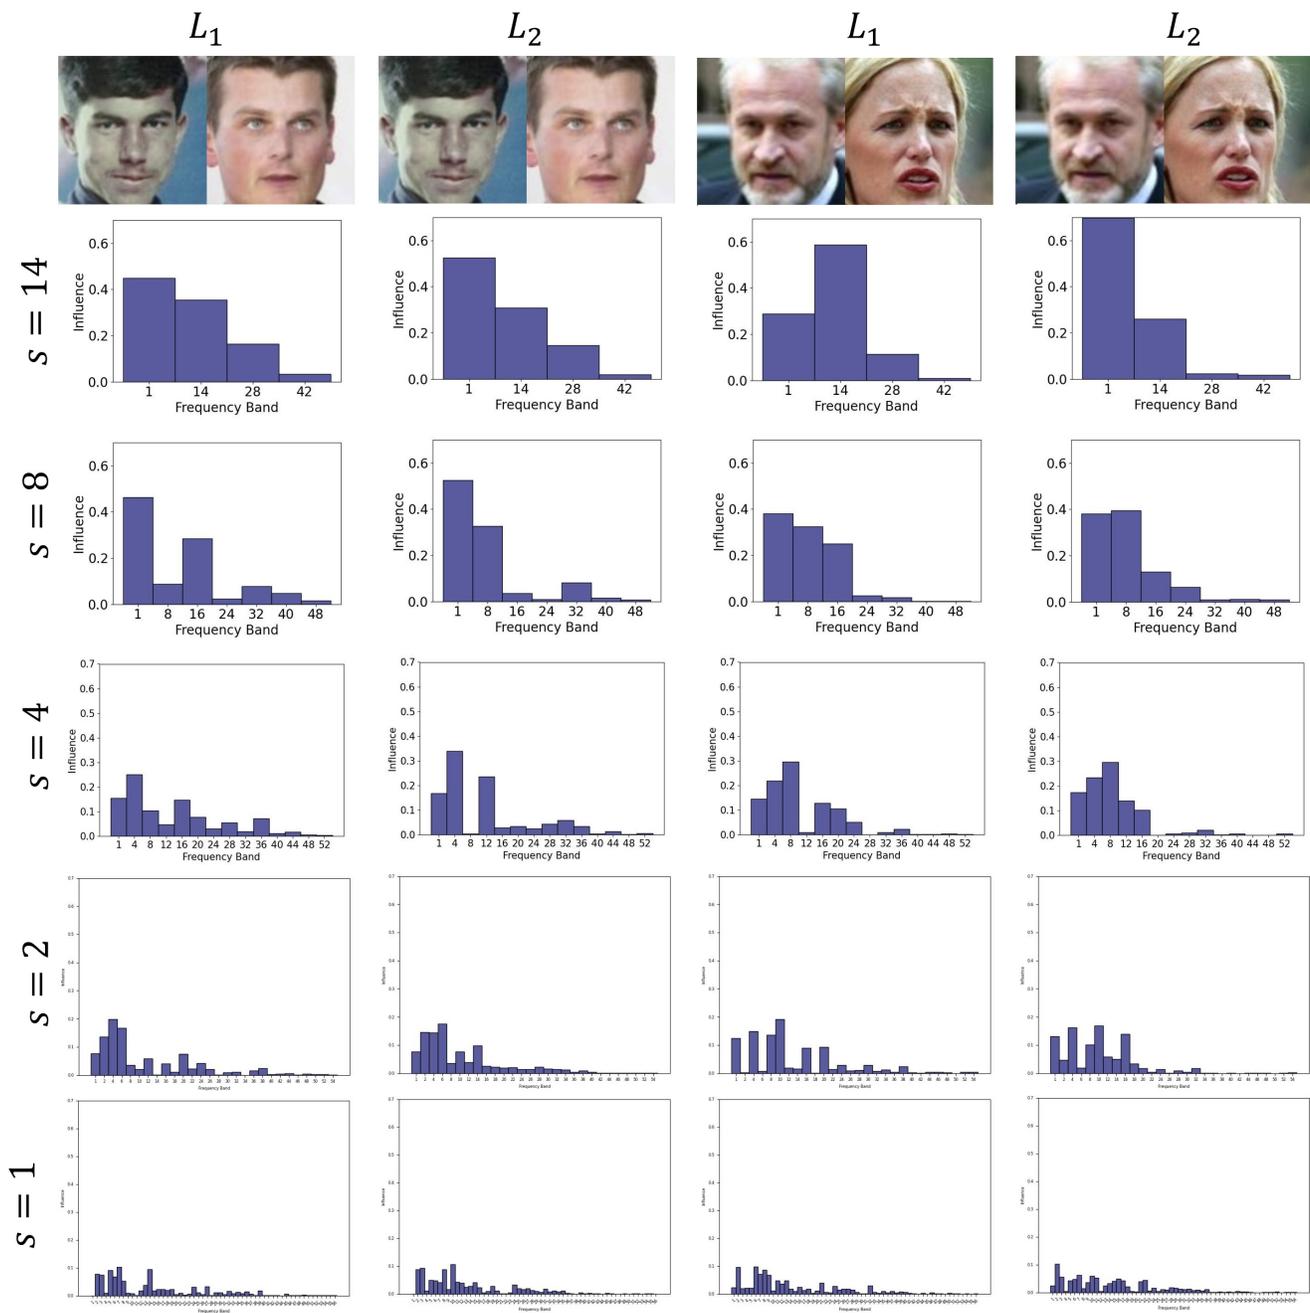


Figure 10. Absolute FHPs using the ElasticFace-Arc [4] model on low-resolution non-matching pairs and either  $\mathcal{L}_1$ - or  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show some similarity independent of the norm utilized for masking, especially with smaller bandsizes. For the sake of comparison, we keep the y-axis scale fixed over all absolute FHPs.

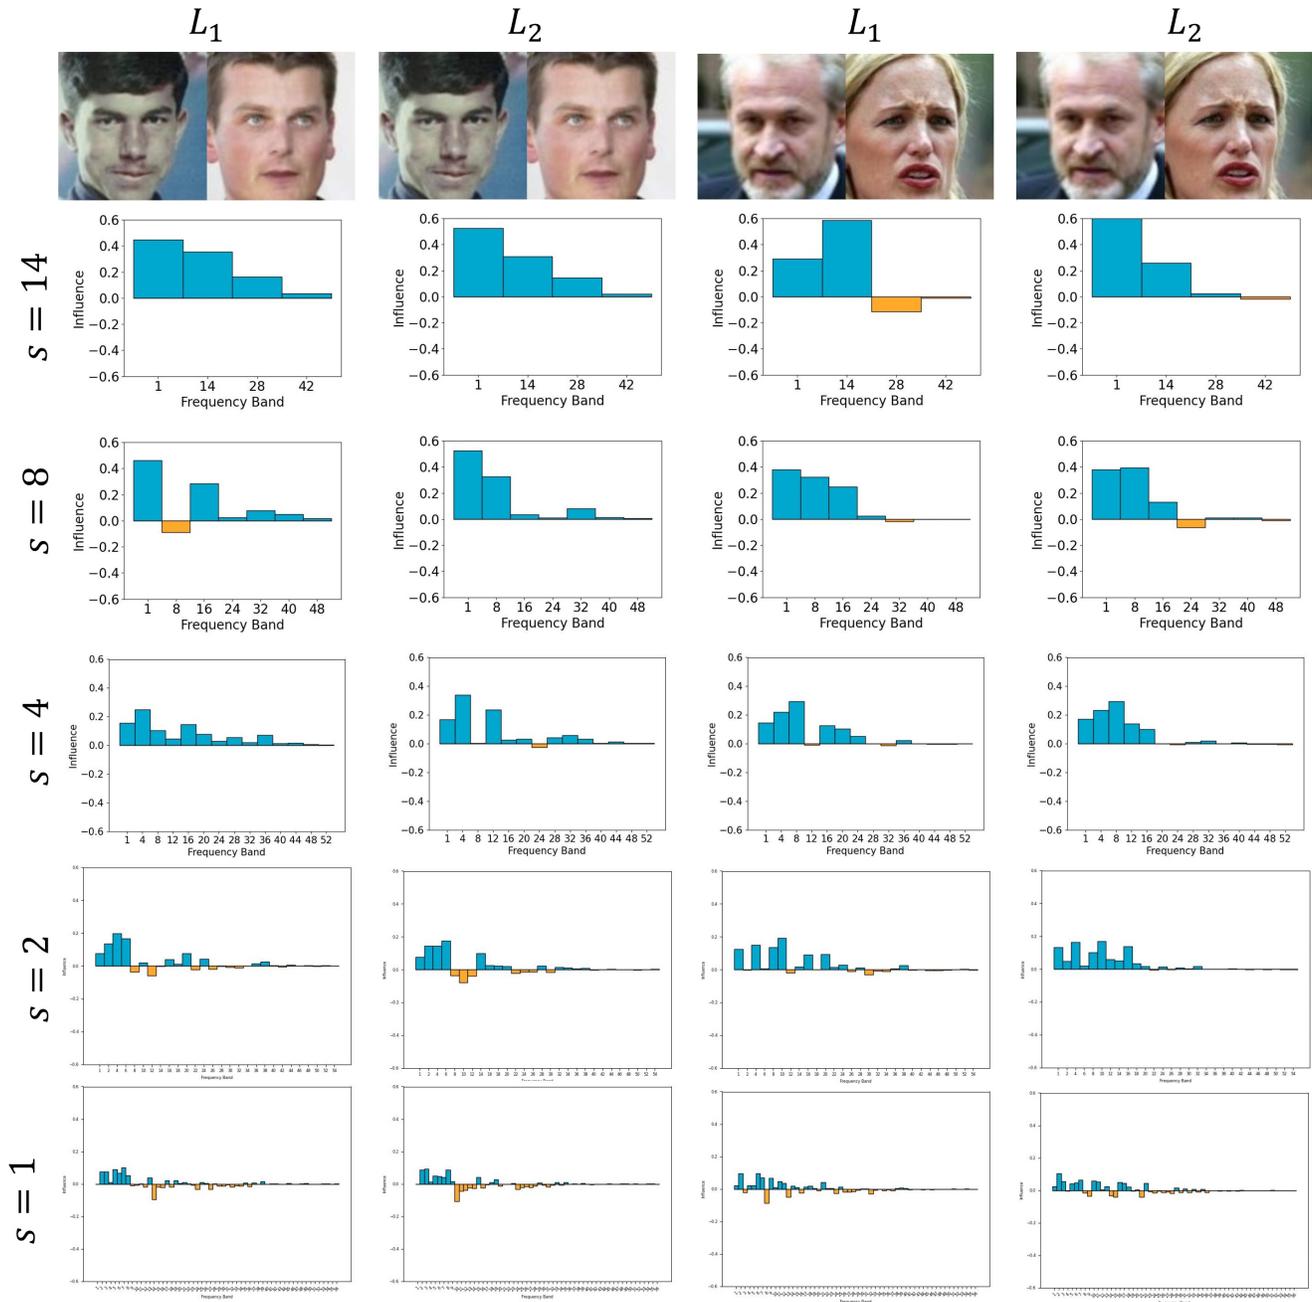


Figure 11. Directed FHPs using the ElasticFace-Arc [4] model on low-resolution non-matching pairs and either  $\mathcal{L}_1$ - or  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show some similarity independent of the norm utilized for masking, especially with smaller bandsizes. For the sake of comparison, we keep the y-axis scale fixed over all directed FHPs.

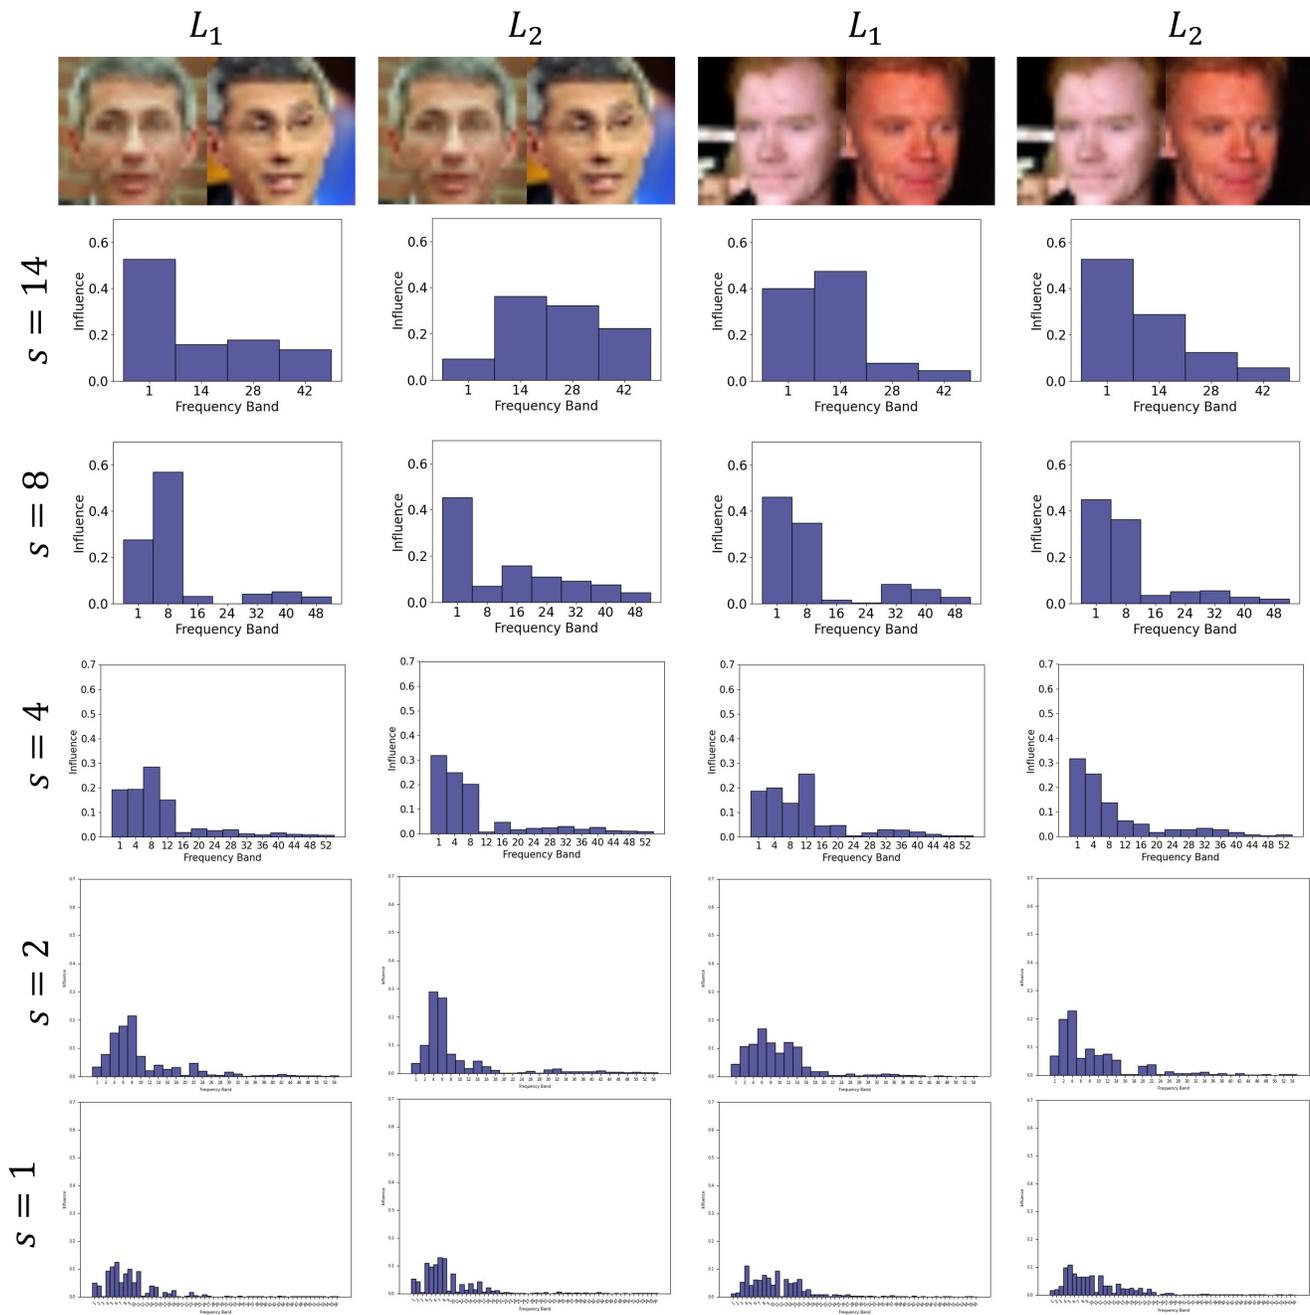


Figure 12. Absolute FHPs using the ElasticFace-Arc [4] model on low-resolution matching pairs and either  $\mathcal{L}_1$ - or  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show some similarity independent of the norm utilized for masking, especially with smaller bandsizes. For the sake of comparison, we keep the y-axis scale fixed over all absolute FHPs.

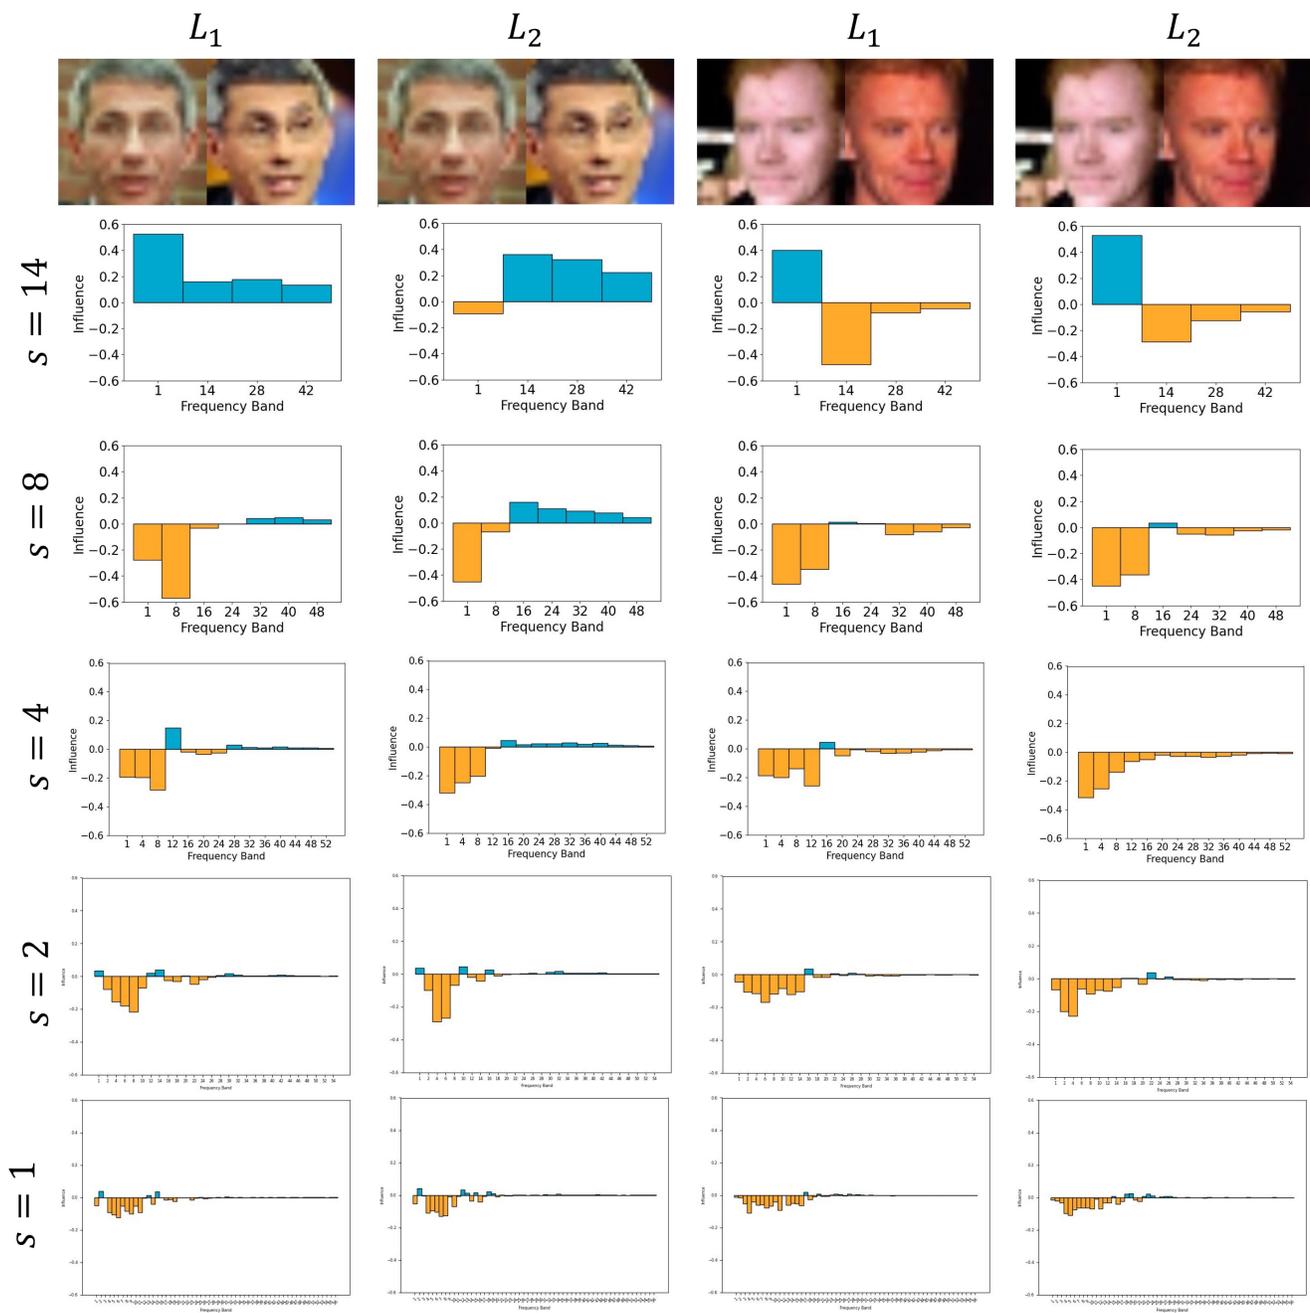


Figure 13. Directed FHPs using the ElasticFace-Arc [4] model on low-resolution matching pairs and either  $\mathcal{L}_1$ - or  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show some similarity independent of the norm utilized for masking, especially with smaller band sizes. For the sake of comparison, we keep the y-axis scale fixed over all directed FHPs.

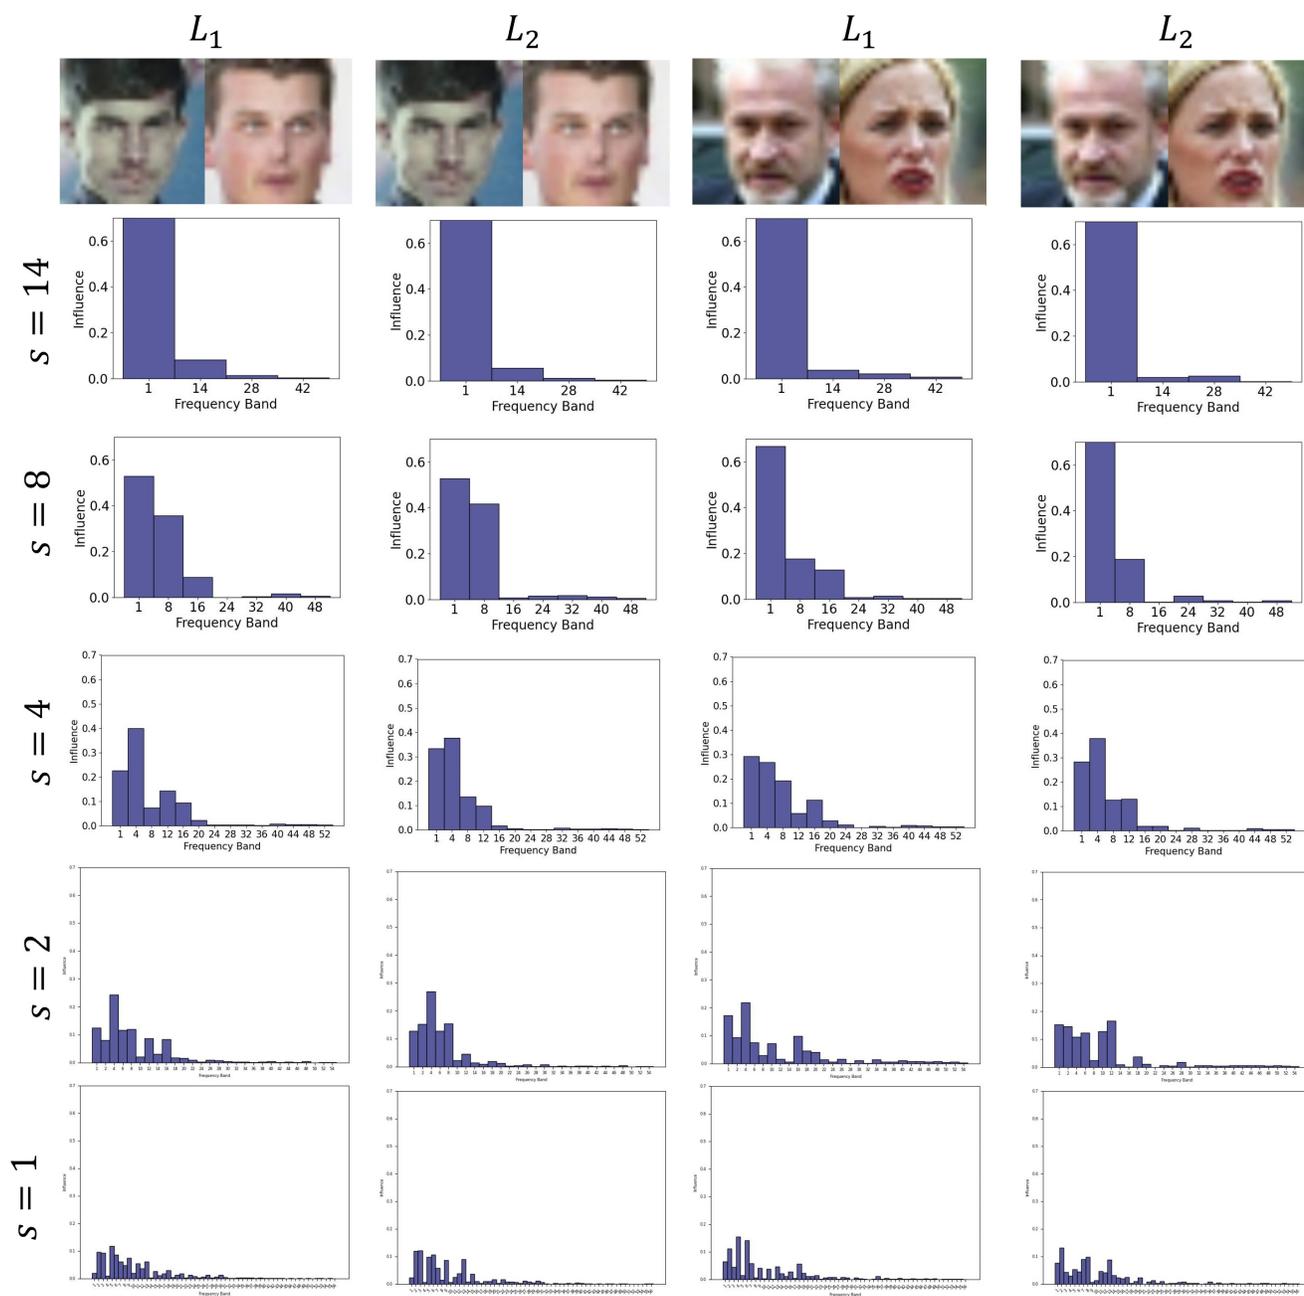


Figure 14. Absolute FHPs using the ElasticFace-Arc [4] model on low-resolution non-matching pairs and either  $\mathcal{L}_1$ - or  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show some similarity independent of the norm utilized for masking, especially with smaller bandsizes. For the sake of comparison, we keep the y-axis scale fixed over all absolute FHPs.

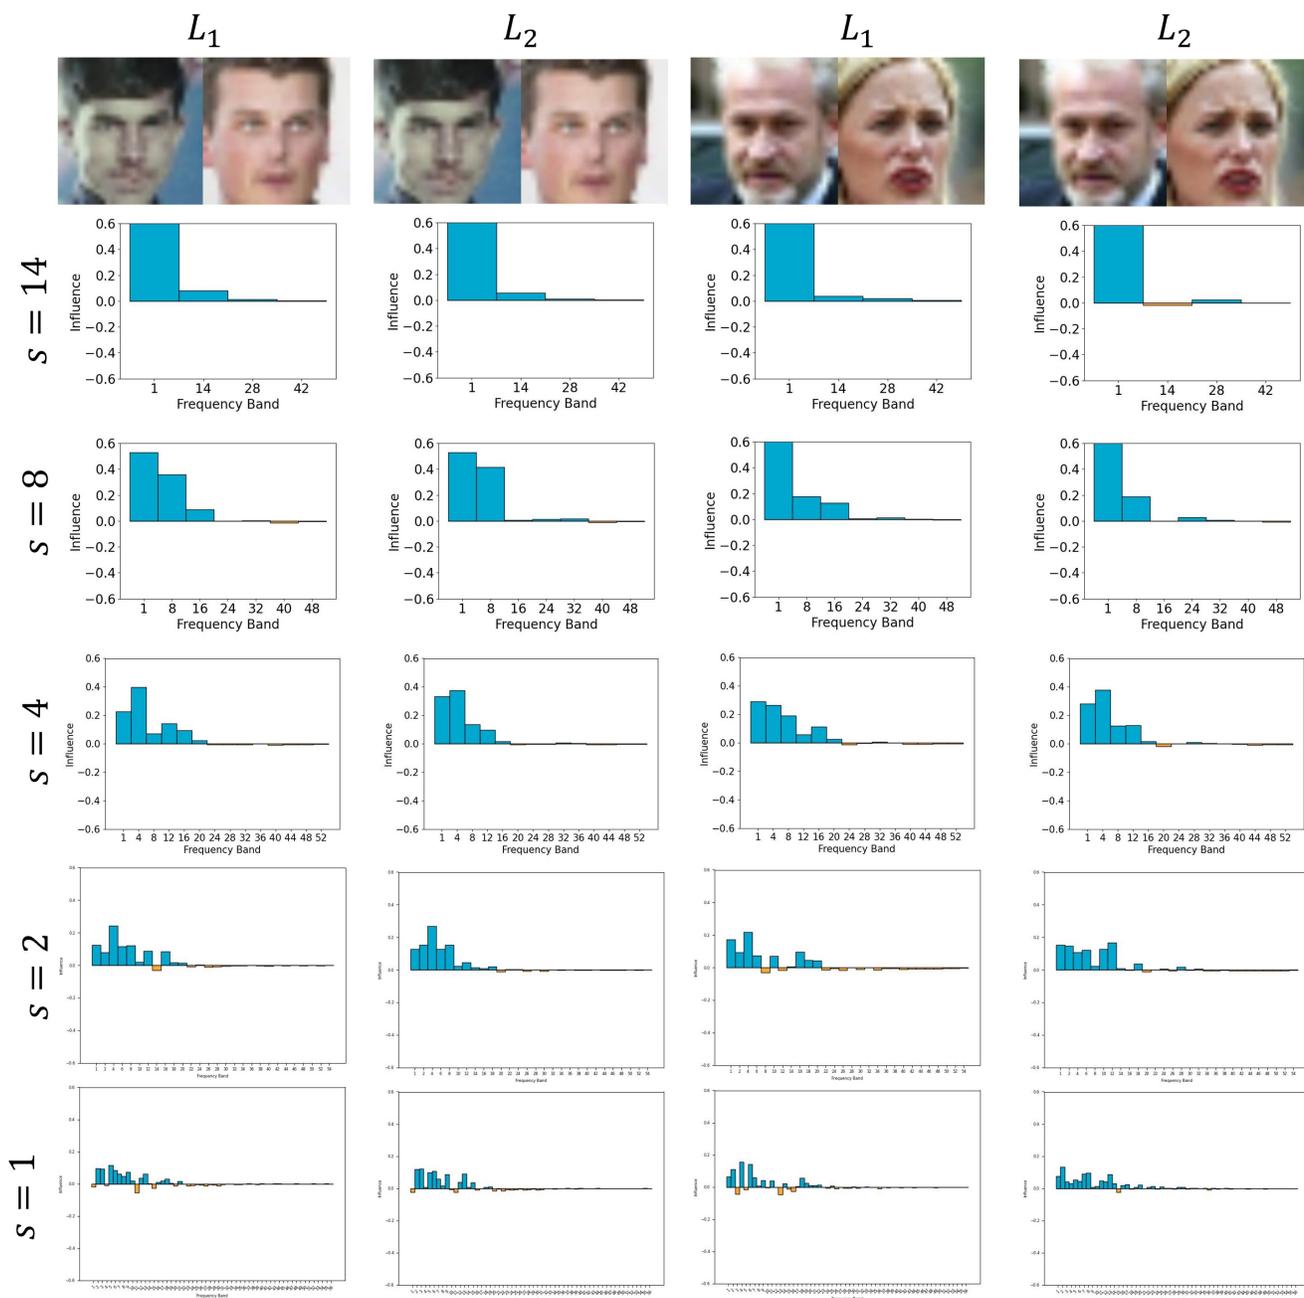


Figure 15. Directed FHPs using the ElasticFace-Arc [4] model on low-resolution non-matching pairs and either  $\mathcal{L}_1$ - or  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show some similarity independent of the norm utilized for masking, especially with smaller bandsizes. For the sake of comparison, we keep the y-axis scale fixed over all directed FHPs.

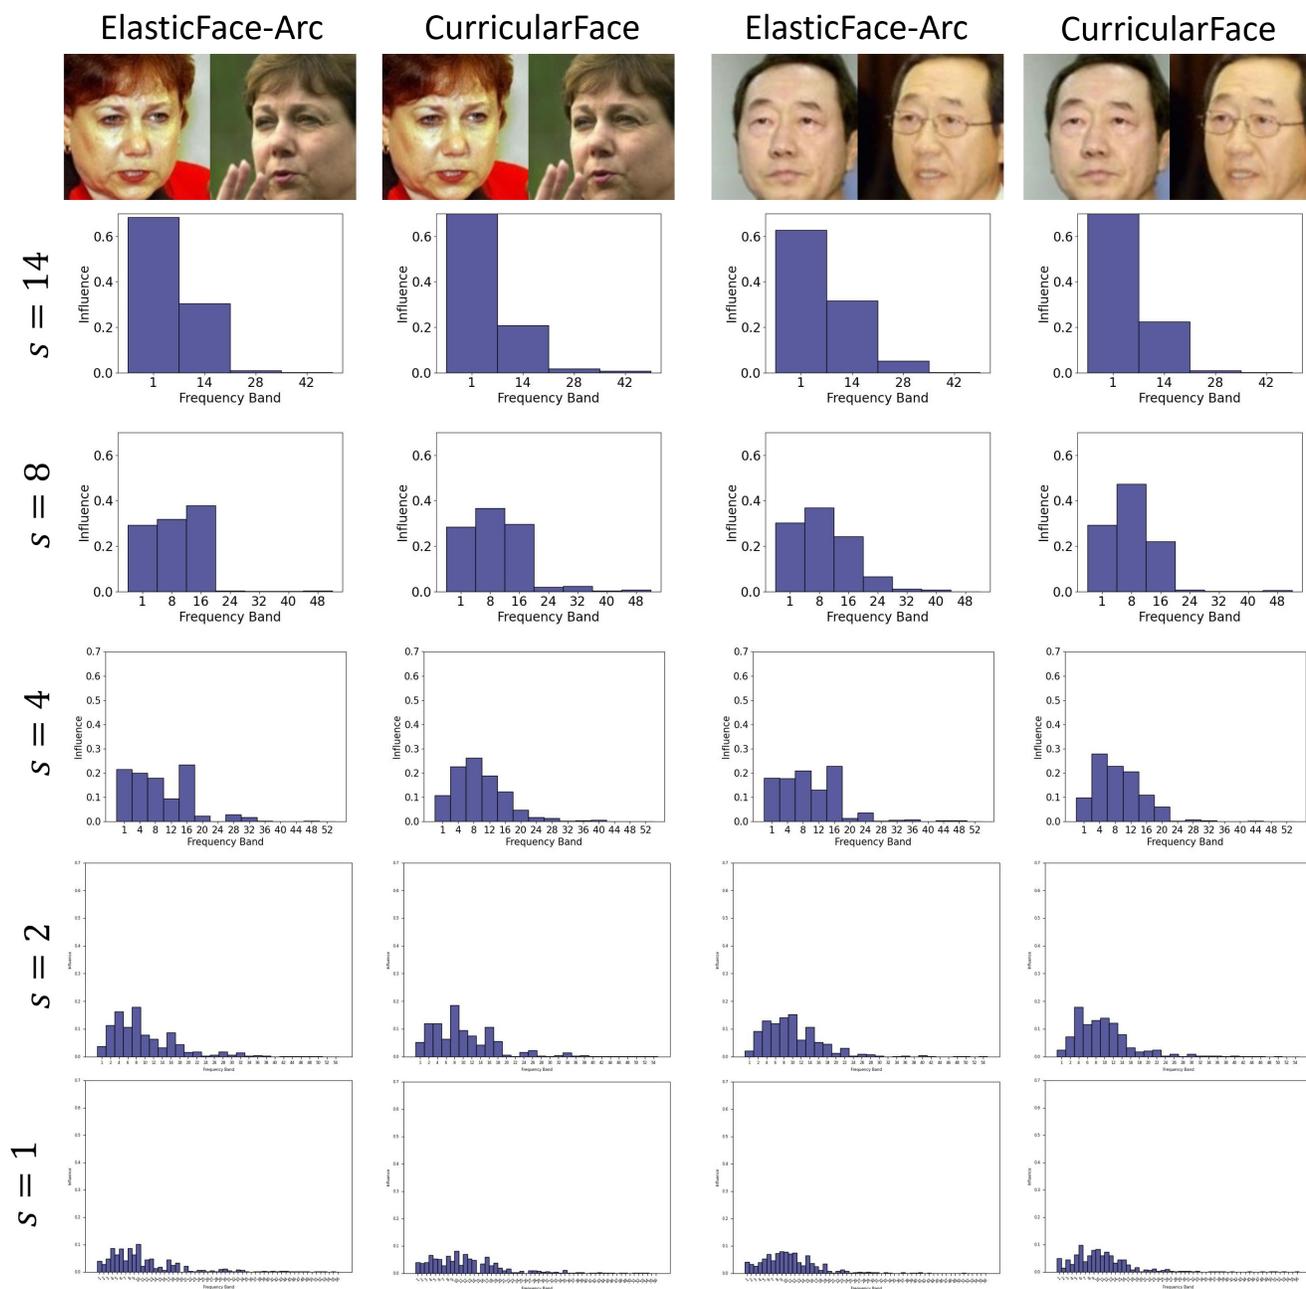


Figure 16. Comparison of absolute FHPs using the ElasticFace-Arc [4] model and the CurricularFace [15] model on unaltered matching pairs using  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show that, to some extent, the same frequency bands have a similar influence independent from the utilized model. For the sake of comparison, we keep the y-axis scale fixed over all absolute FHPs.

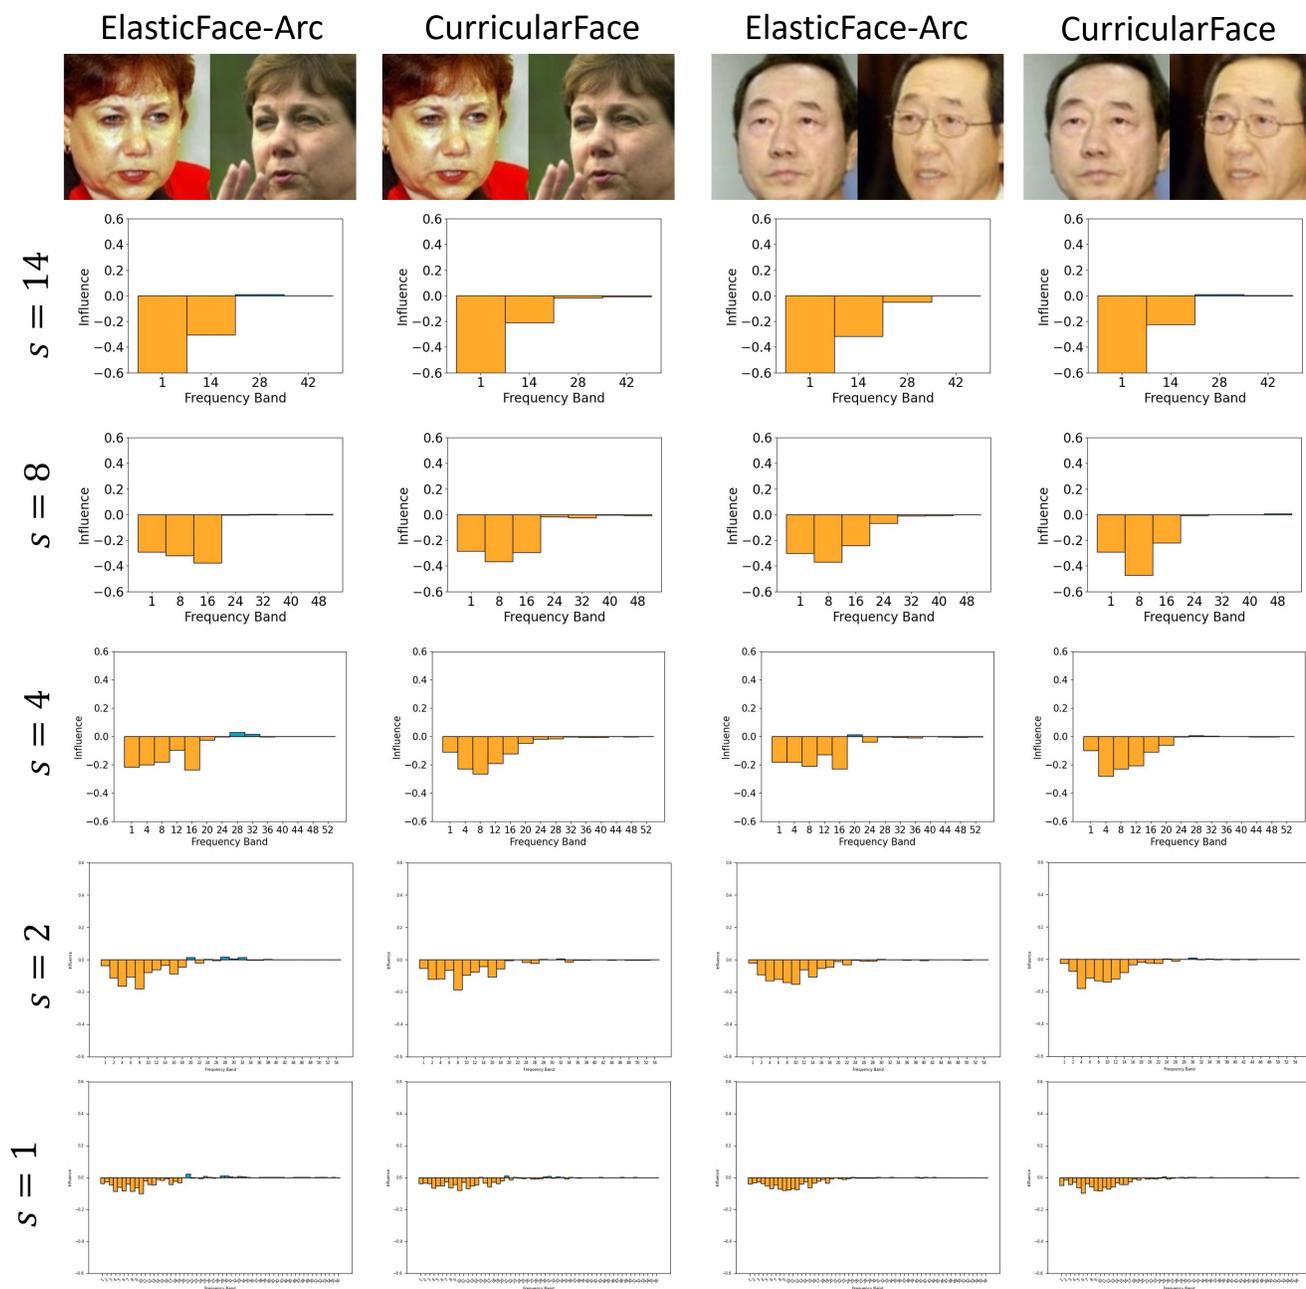


Figure 17. Comparison of directed FHPs using the ElasticFace-Arc [4] model and the CurricularFace [15] model on unaltered matching pairs using  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show that, to some extent, the same frequency bands have a similar influence independent from the utilized model. For the sake of comparison, we keep the y-axis scale fixed over all directed FHPs.

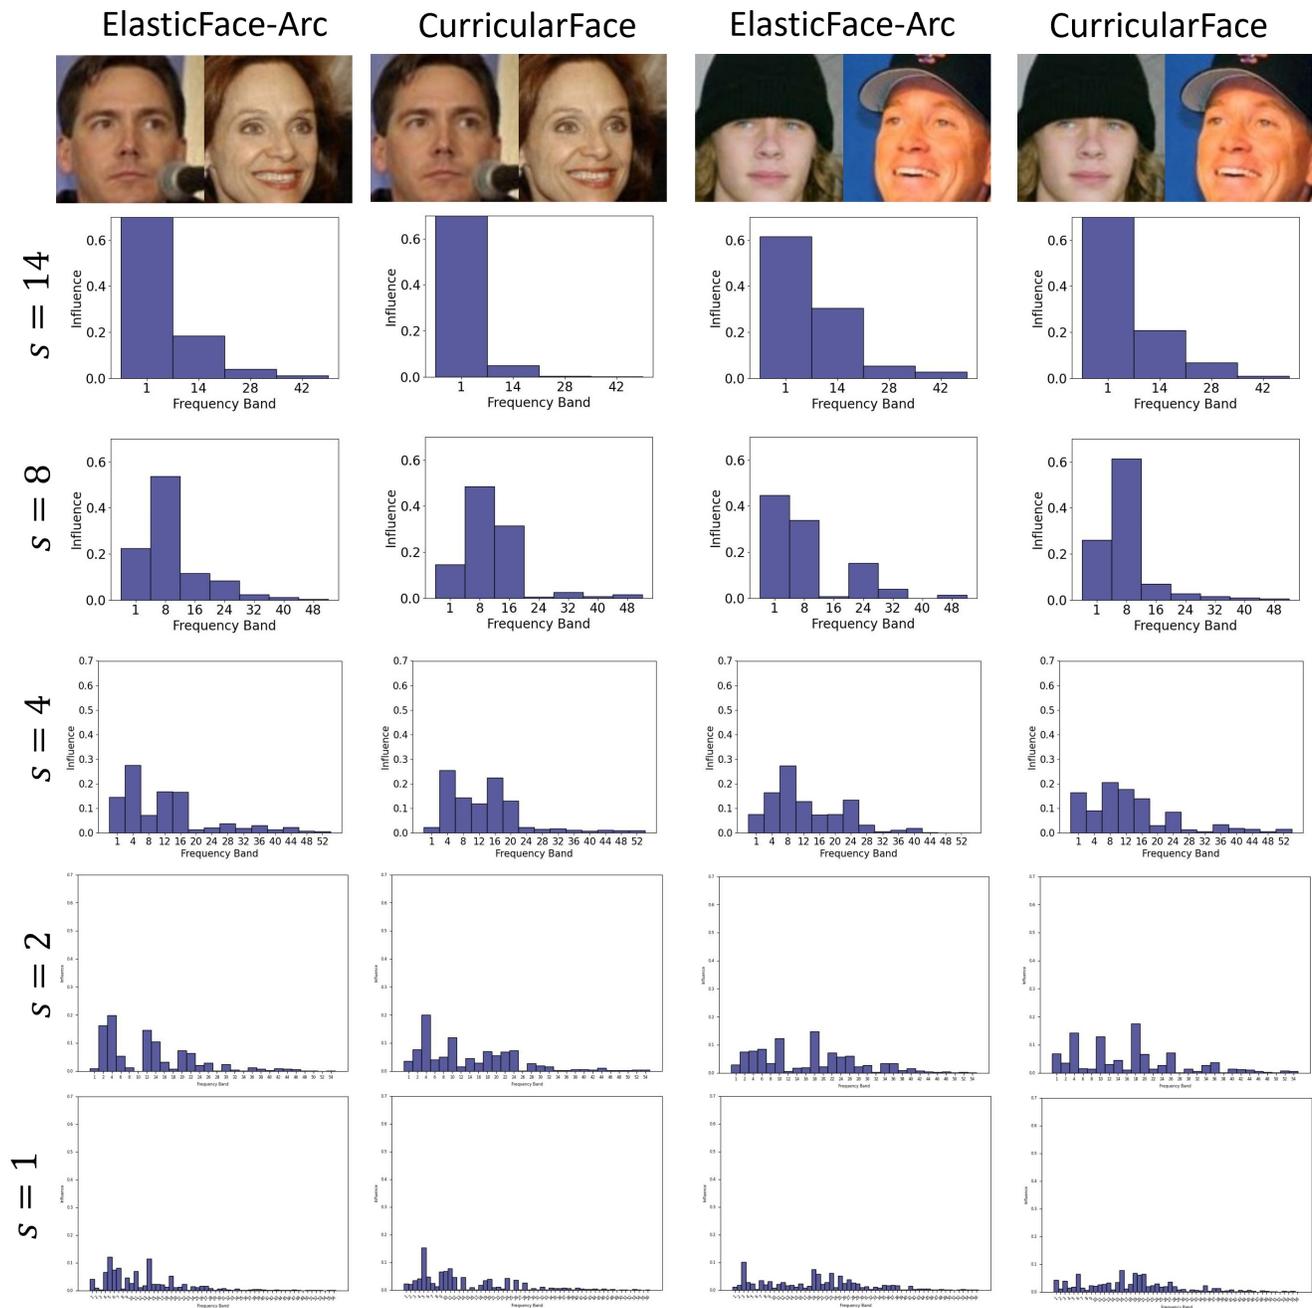


Figure 18. Comparison of absolute FHPs using the ElasticFace-Arc [4] model and the CurricularFace [15] model on unaltered non-matching pairs using  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show that, to some extent, the same frequency bands have a similar influence independent from the utilized model. For the sake of comparison, we keep the y-axis scale fixed over all absolute FHPs.

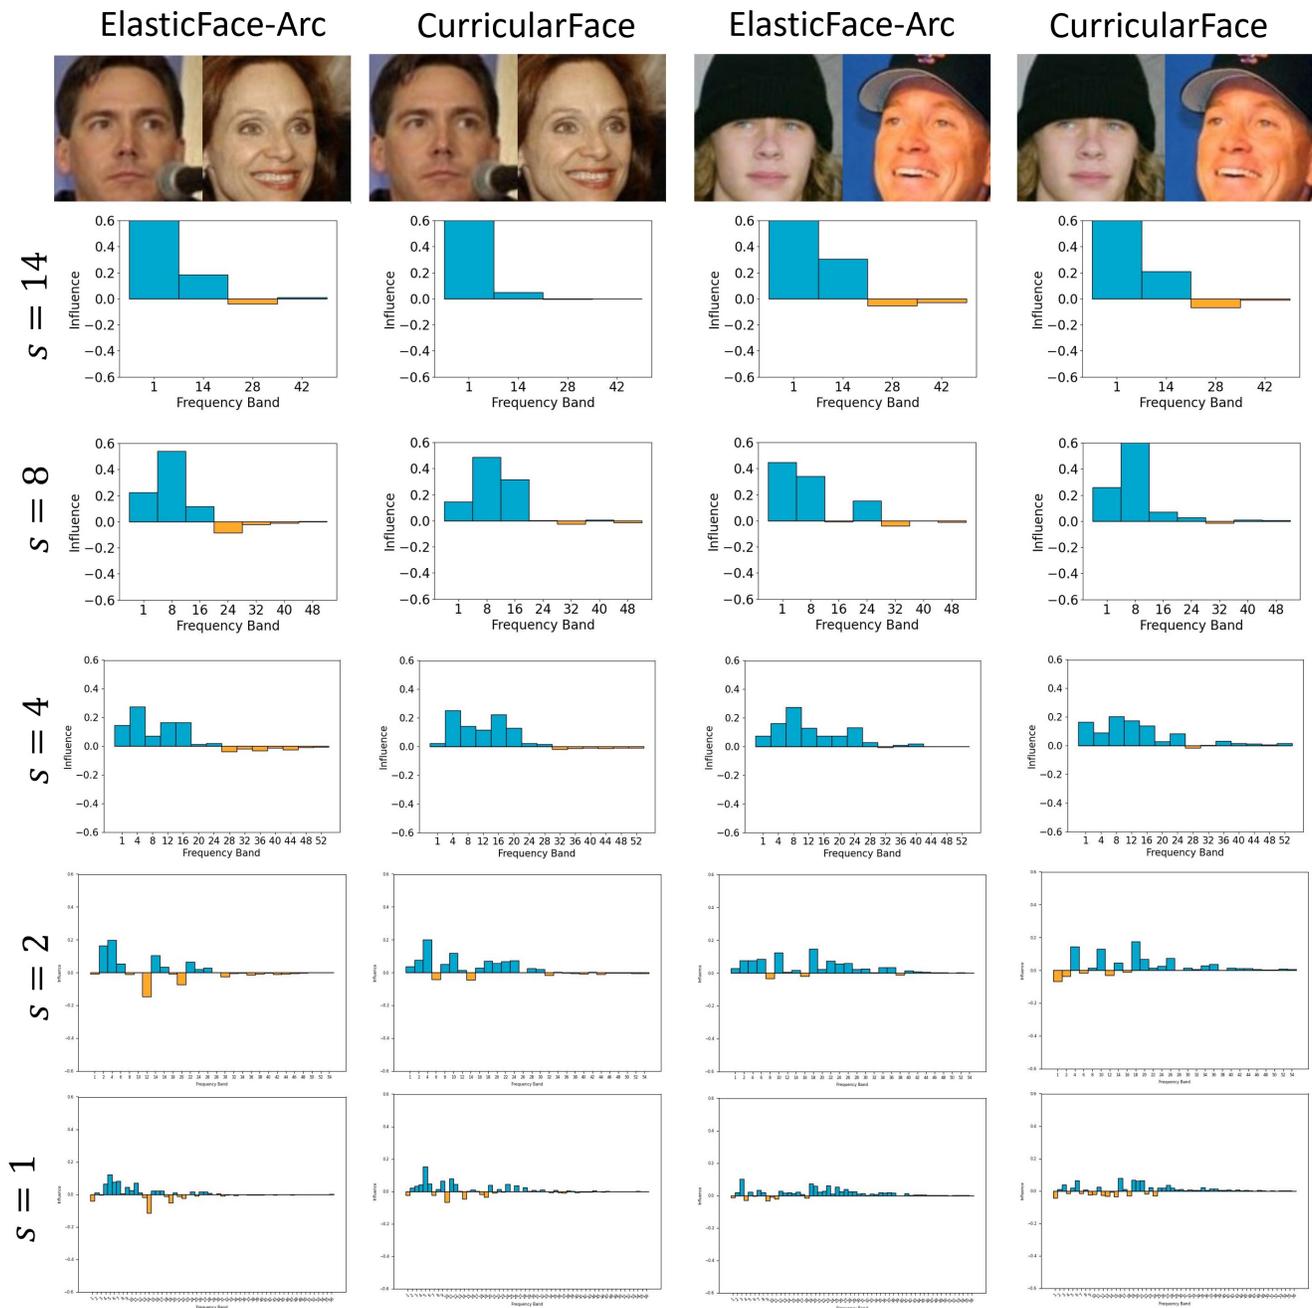


Figure 19. Comparison of directed FHPs using the ElasticFace-Arc [4] model and the CurricularFace [15] model on unaltered non-matching pairs using  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show that, to some extent, the same frequency bands have a similar influence independent from the utilized model. For the sake of comparison, we keep the y-axis scale fixed over all directed FHPs.

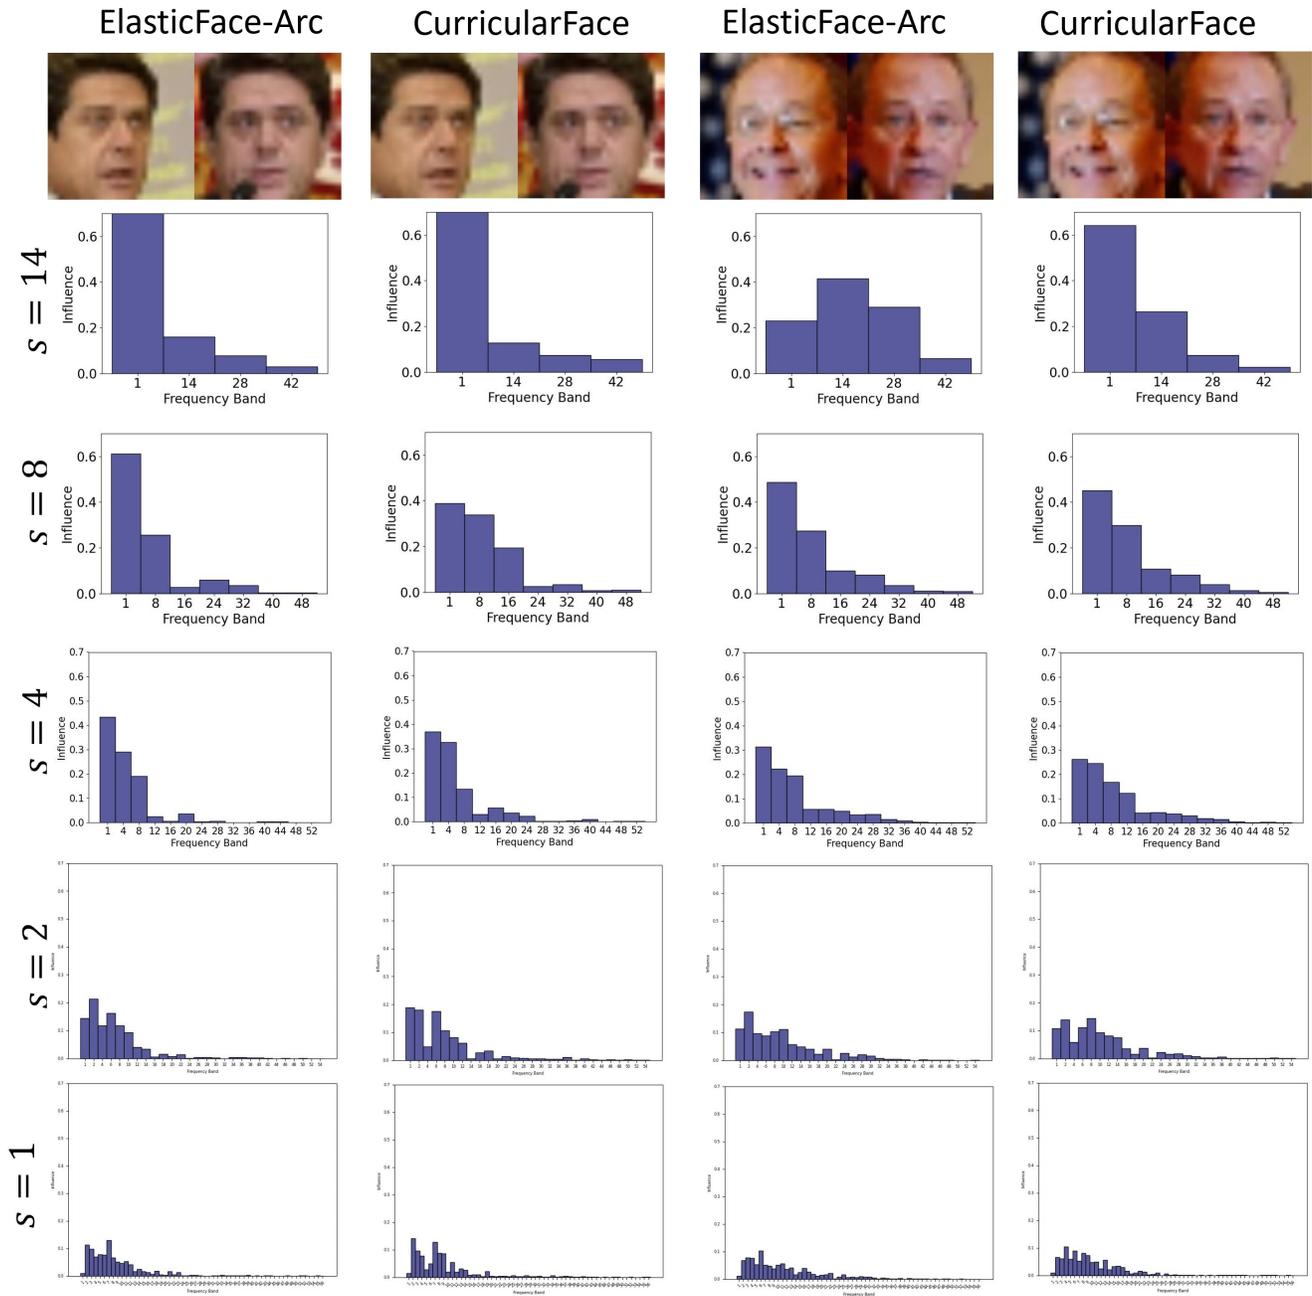


Figure 20. Comparison of absolute FHPs using the ElasticFace-Arc [4] model and the CurricularFace [15] model on low-resolution matching pairs using  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show that, to some extent, the same frequency bands have a similar influence independent from the utilized model. For the sake of comparison, we keep the y-axis scale fixed over all absolute FHPs.

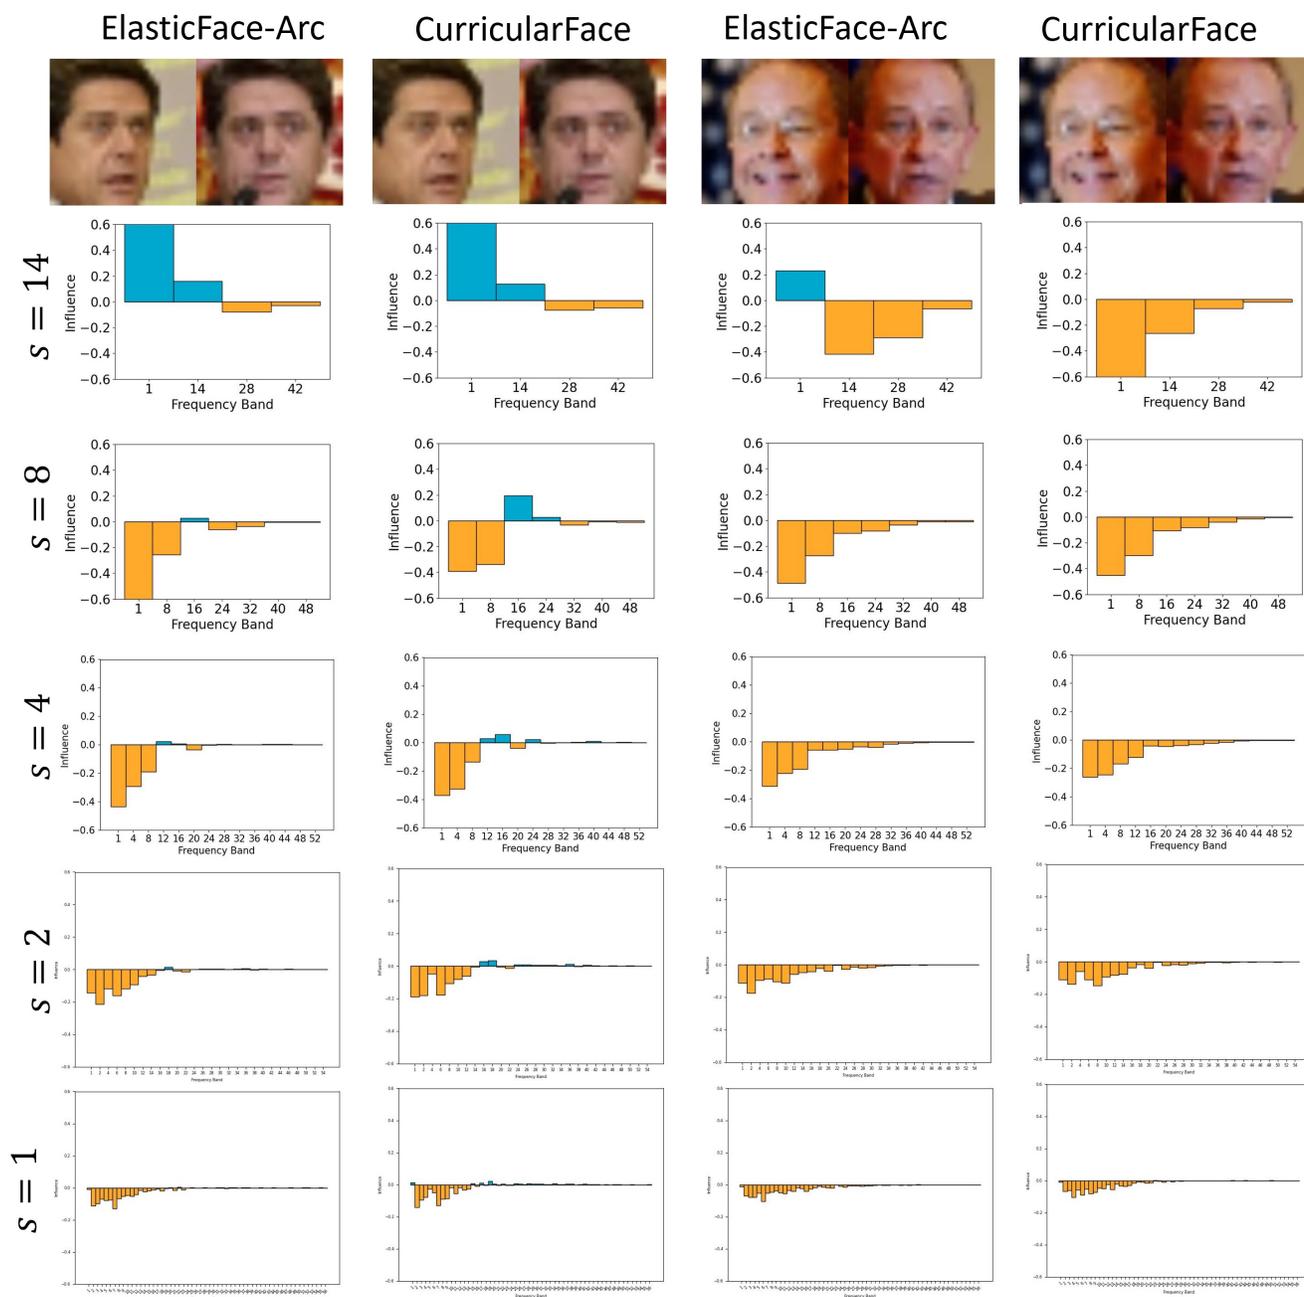


Figure 21. Comparison of directed FHPs using the ElasticFace-Arc [4] model and the CurricularFace [15] model on low-resolution matching pairs using  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show that, to some extent, the same frequency bands have a similar influence independent from the utilized model. For the sake of comparison, we keep the y-axis scale fixed over all directed FHPs.

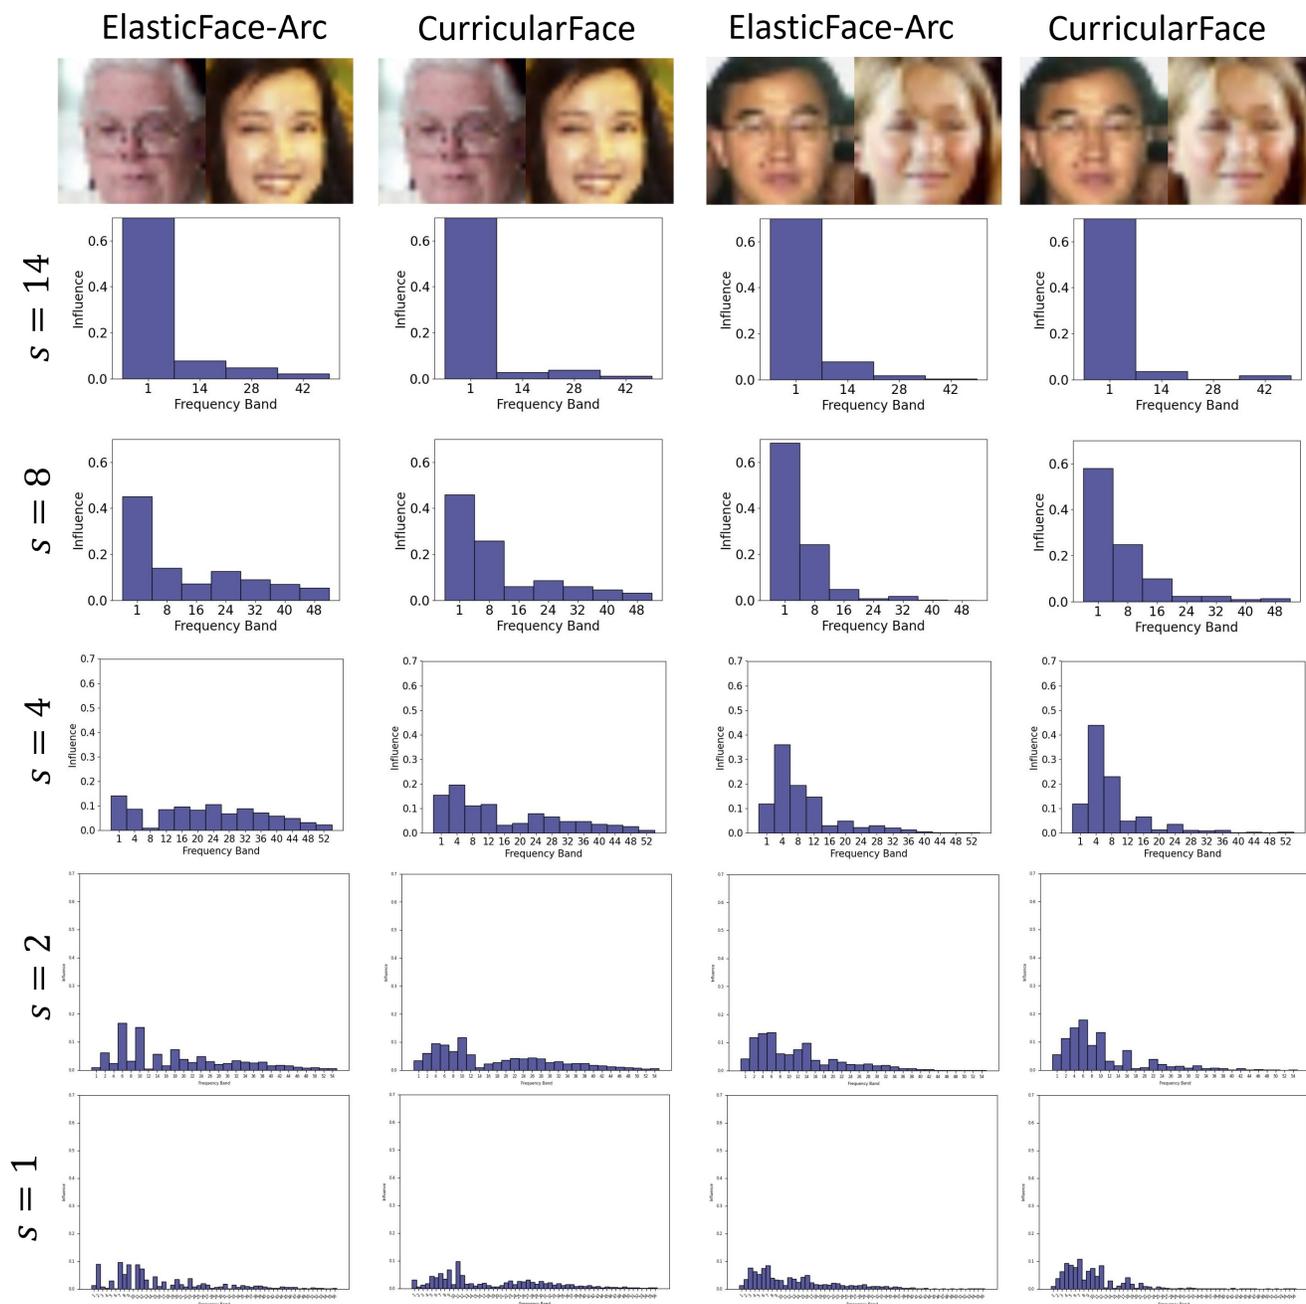


Figure 22. Comparison of absolute FHPs using the ElasticFace-Arc [4] model and the CurricularFace [15] model on low-resolution non-matching pairs using  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show that, to some extent, the same frequency bands have a similar influence independent from the utilized model. For the sake of comparison, we keep the y-axis scale fixed over all absolute FHPs.

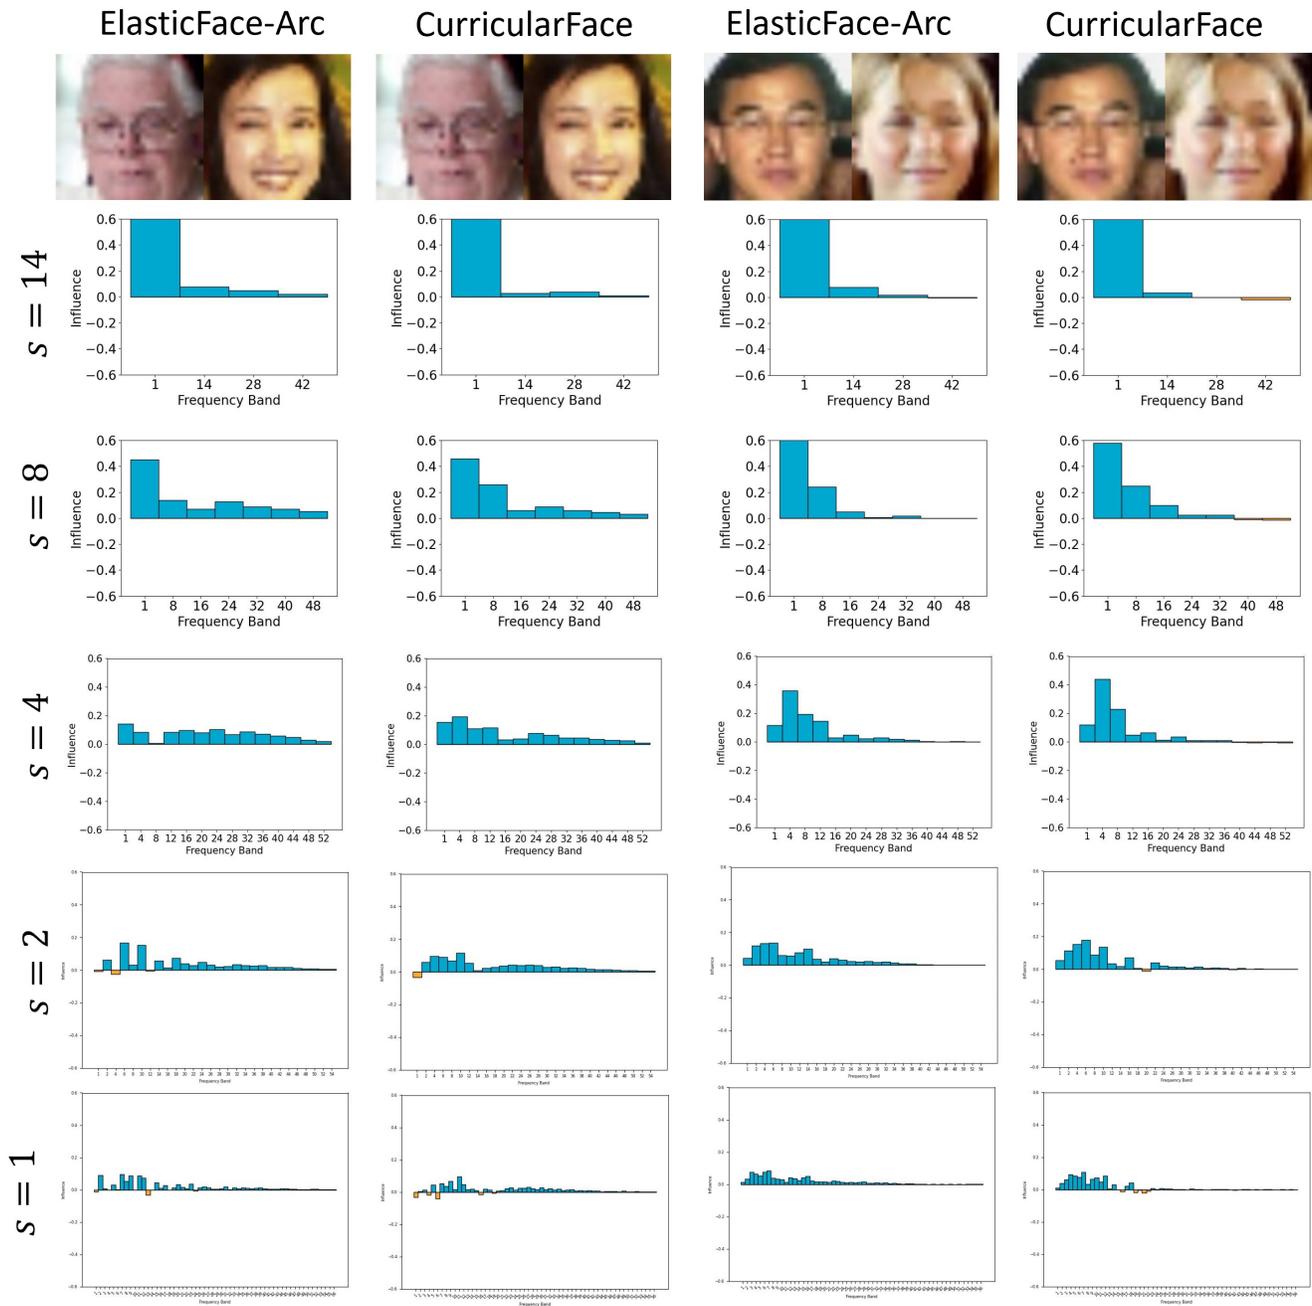


Figure 23. Comparison of directed FHPs using the ElasticFace-Arc [4] model and the CurricularFace [15] model on low-resolution non-matching pairs using  $\mathcal{L}_2$ -norm during the masking process. The distribution of the influences show that, to some extent, the same frequency bands have a similar influence independent from the utilized model. For the sake of comparison, we keep the y-axis scale fixed over all directed FHPs.