EFFECT OF LOSSY COMPRESSION ALGORITHMS ON FACE IMAGE QUALITY AND RECOGNITION

Torsten Schlett   *Sebastian Schachner*   *Christian Rathgeb*   *Juan Tapia*   *Christoph Busch*

Hochschule Darmstadt, Germany

ABSTRACT

Lossy face image compression can degrade the image quality and the utility for the purpose of face recognition. This work investigates the effect of lossy image compression on a state-of-the-art face recognition model, and on multiple face image quality assessment models. The analysis is conducted over a range of specific image target sizes. Four compression types are considered, namely JPEG, JPEG 2000, downscaled PNG, and notably the new JPEG XL format. Frontal color images from the ColorFERET database were used in a Region Of Interest (ROI) variant and a portrait variant. We primarily conclude that JPEG XL allows for superior mean and worst case face recognition performance especially at lower target sizes, below approximately 5kB for the ROI variant, while there appears to be no critical advantage among the compression types at higher target sizes. Quality assessments from modern models correlate well overall with the compression effect on face recognition performance.

Index Terms— Biometrics, face recognition, face image quality assessment, lossy image compression

1. INTRODUCTION

Image quality in terms of biometric utility [1] influences the performance of face recognition systems, as e.g. used for automated border control. Lossy image compression can affect the utility, and is relevant e.g. for Machine Readable Travel Documents (MRTDs) [2]. Since the effect on utility may differ substantially between lossy compression schemes at approximately the same compressed image size, this leads to the question which compression type should be preferred. Furthermore, there are automatic Face Image Quality Assessment (FIQA) methods that attempt to estimate an individual image's biometric utility, leading to the question whether these
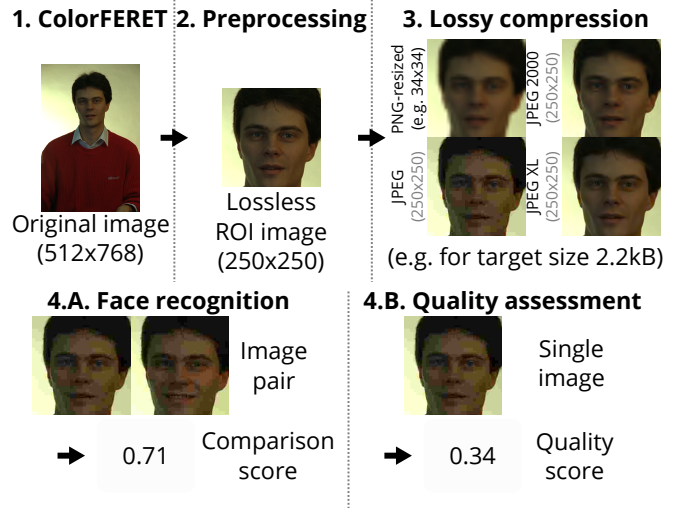
Fig. 1. Processing stages required for the experiments.

methods' assessments are similarly influenced by image compression in isolation.

As an example for related prior work, Funk et al. [3] evaluated JPEG and JPEG 2000 with respect to two unspecified face recognition systems on 358 frontal face images from the FERET [4][5] database, all 512×768 pixels, finding that JPEG 2000 outperformed JPEG below a size of 17kB. Later work by Delac et al. [6] likewise evaluated JPEG and JPEG 2000 on images from FERET, but for different grayscale subsets transformed and cropped to a 128×128 pixels ROI with masked background, for three face recognition algorithms (based on Principal Component Analysis, Linear Discriminant Analysis, Independent Component Analysis), similarly finding JPEG 2000 results superior to JPEG for higher compression rates. Our work contributes an evaluation of the effect of lossy compression on face recognition using a modern ArcFace [7] model, comparing the new JPEG XL [8] encoding besides the previously evaluated JPEG [9] and JPEG 2000 [10], as well as downscaled PNG [11] images, which proved to be competitive with the dedicated lossy compression methods. Our evaluation makes use of frontal color images from ColorFERET [4][5] across a range of concrete image target sizes. Additionally, we examined how well the effect of lossy compression on face recognition is reflected in the per-image quality scores produced by eight

different FIQA methods, which previously have only rarely been examined in a lossy image compression context [12]. Figure 1 provides an overview of the processes required for the experiment data.

In the remainder of this paper, section 2 describes the details of the experiment setup, section 3 presents an analysis of the results, and section 4 concludes with the main findings.

2. SETUP

2.1. Database

A high initial quality of the images is desired for the experiments, to better isolate the impact of lossy image compression for the analysis. Focusing on frontal images further allows us to approximate a passport image scenario. Thus the frontal color images from the ColorFERET [4][5] database were used for the experiments. This amounts to 2638 images across 966 subjects. These images were captured under comparatively controlled conditions, in contrast to e.g. a web-scraped face image database such as LFW [13]. The facial expression does however vary (neutral or smiling), and for some subjects there are images both with and without glasses.

Many of the original ColorFERET [4][5] images show a larger area than just the face, so we applied preprocessing to obtain ROI images for the experiments, as illustrated in Figure 1. Images were cropped and aligned as required for direct usage by the ArcFace [7] face recognition model, except for the resolution, which was 250×250. The landmark detection required for preprocessing was done via the publicly available RetinaFace-R50 model [14] from the Insight-Face project. Both the face recognition model and the six deep learning FIQA models require various fixed input image sizes, all smaller than 250×250, thus images were resized with bilinear interpolation for each of these models.

2.2. Face recognition

To obtain comparison scores [1] we used the publicly available ArcFace-R100-MS1MV2 [7] face recognition model from the InsightFace project, which requires a 112×112 input image size. Multiple comparison trial [1] configurations were analyzed: The "mated-other" configuration consisted out of all 3484 possible mated [1] trials for the subjects, and compared different captured images from the same subject with the same target size compression applied. The "mated-self" configuration on the other hand consisted out of comparisons of the 2638 images with lossless compression against the same face image (stemming from the same captured sample) with lossy compression applied.

2.3. Quality assessment

To obtain quality scores we used eight FIQA methods. For all of the methods, higher quality score output is intended to indicate higher biometric utility [1]. Six of these were modern deep learning models specifically intended to assess quality in terms of image utility [1] for face recognition: CR-FIQA(S) & CR-FIQA(L) [15] (respectively

with iResNet50/iResNet100 backbone trained on CASIA-WebFace[16]/MS1MV2[7], both with 112×112 input image size), MagFace [17] (iResNet100 backbone trained on MS1MV2 [7], 112×112 input image size), SER-FIQ [18] ("same model" variant using ArcFace, 112×112 input image size), FaceQnet-v0 [19] & FaceQnet-v1 [20] (both with ResNet-50 backbone, trained on VGGFace2 [21], 224×224 input image size). These six models are all publicly available. The other two FIQA methods were handcrafted sharpness measures, referred to as Sharpness-1 and -2: Sharpness-1 is implemented as described by [22], comparing the pixel intensity variations in the input image against those in an artificially blurred version thereof to estimate the sharpness. Sharpness-2 is conceptually similar but simpler, computing the sharpness as the mean of the distances between the normalized $[0, 1]$ pixel values of the input image and a blurred (3×3 mean box filter) version thereof.

2.4. Compression target sizes and types

Target sizes for the lossy image compression were selected manually, based on the lowest usable size among all used compression types, and based on the sizes of the images as lossless PNG files. Images were compressed as closely as possible to each target size, with the target size serving as the inclusive upper limit, meaning that an image's size may at most be equal to the target size. We selected the target sizes 5kB, 4.5kB, 4kB, 3.5kB, 3kB, 2.5kB, and 2.2kB.

Four lossy compression types were examined: PNG-resized, JPEG [9], JPEG 2000 [10], and JPEG XL [8]. PNG-resized refers to PNG [11] images that were downscaled to reach each target size, keeping the aspect ratio of the source image as closely as possible, thus effectively turning the by itself lossless PNG compression into a lossy compression scheme. PNG, JPEG and JPEG 2000 are permissible as face image encodings in MRTDs according to ISO/IEC 39794-5:2019 [2]. JPEG XL is a more recently created and standardized encoding. For PNG, JPEG, and JPEG 2000 compression we used Pillow version 9.1.0, which in turn relies on ZLIB version 1.2.12 for PNG, on libjpeg-turbo version 2.1.3 for JPEG, and on OpenJPEG version 2.4.0 for JPEG 2000. For JPEG XL we used the reference implementation libjxl version 0.6.1. Default compression settings were used for all compression types, except for the comparatively unusual PNG-resized compression, for which we did set the Pillow "optimize" option.

Step three in Figure 1 shows examples for ROI images compressed with all four types at target size 2.2kB.

3. RESULTS

3.1. Effect on comparison scores

The comparison scores shown in Figure 2 are similarity scores [1] in the range $[-1, +1]$, meaning that higher values are intended to indicate higher facial similarity. The y-axis shows the comparison score, and the x-axis shows the target sizes in
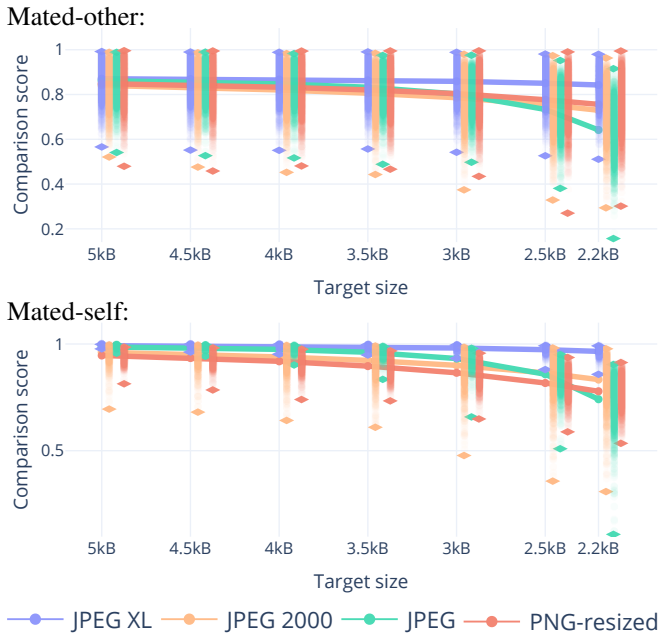
Fig. 2. Comparison score results.

Fig. 3. Normalized mean quality and mated-self/other comparison score curves.

descending order from left to right. The curves correspond to the different compression types. Each linearly interpolated curve shows the mean value of the comparison scores per target size. The transparent and thus "smeared" looking markers of the same color show all concrete individual comparison score data points per target size, whereby the x-axis location is slightly shifted across the compression types for the sake of visibility. Additional opaque rhombus-shaped markers highlight the maximum and minimum comparison scores.

In the results for the mated-other comparison trial configuration, JPEG XL performed best both in terms of maintaining high mean comparison scores and in terms of the worst-case (minimum) comparison scores. The compression types' results diverged most distinctly for the lowest target size, 2.2kB. At this target size, JPEG's performance in terms of the mean score and in terms of the overall worst-case score dropped noticeably, relative to slightly higher target sizes at which JPEG's results remained closer to or better than JPEG 2000 and PNG-resized. A perhaps surprising result can be seen for PNG-resized, which remained competitive with JPEG and JPEG 2000 across the target sizes, and even yielded slightly better results than both at 2.2kB.

Results for the mated-self comparison trial configuration lead to similar conclusions as for mated-other, with some differences. JPEG XL is again the compression type with the best results, here with worst-case scores that are even more distinguished from those of the other compression types. JPEG is again the worst performing compression at 2.2kB, while being competitive or better than JPEG 2000 and PNG-resized at higher target sizes. Regarding differences to mated-other, the order of the JPEG 2000 and PNG-resized mean scor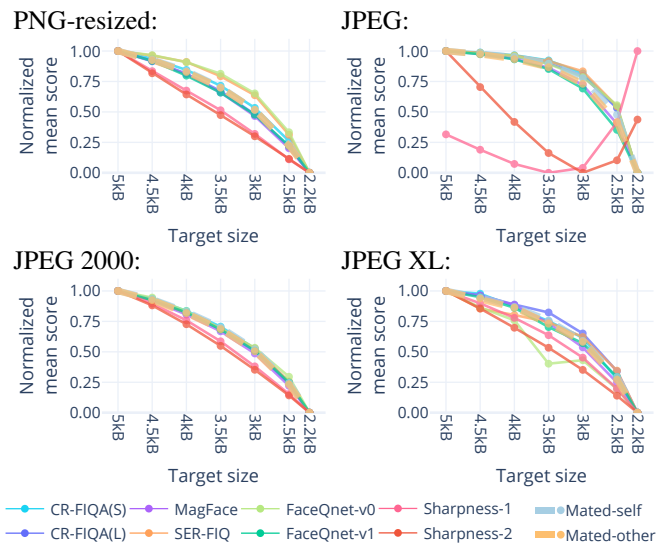e curves is reversed, JPEG 2000 now being slightly better in this metric, although PNG-resized now beats JPEG 2000 in terms of worst-case scores across all target sizes.

To summarize, the results let us recommend JPEG XL across all target sizes, while JPEG should be avoided for very low sizes (2.2kB), but is otherwise competitive with JPEG 2000 and PNG-resized. For the mated-self score spread and worst-case results, the lead of the JPEG XL results was more pronounced even at higher shown target sizes, with only JPEG showing competitive results at approximately 4kB and above.

3.2. Effect on quality scores

A direct computation of the correlation between the quality scores and the comparison scores is not possible, since a quality score corresponds to one individual image, while a comparison score corresponds to an image pair. Thus to analyze the extent to which the effect of the lossy image compression on the comparison scores is reflected in the quality scores of the various tested FIQA methods, we chose to examine the normalized mean score curves as shown in Figure 3. The x-axis shows the same target size range as for the comparison scores in Figure 2, while the y-axis now shows the mean quality or comparison scores, normalized to $[0, 1]$ since the original value ranges differed. The plots are separated by the four compression types, and the various curves represent the eight FIQA methods as well as the comparison scores for the mated-self and mated-other trial configuration. These normalized curves show that the effect of the lossy image compression types on most of the FIQA methods is similar to the effect on the comparison scores. The two sharpness measures are the most clear exception, especially for JPEG, where the severe sharpness value increase for target size 2.5kB and 2.2kB can be explained by the more severe JPEG "block" compression artifacts. To better quan-

		CR-FIQA(S)	CR-FIQA(L)	MagFace	SER-FIQ	FaceQnet-v0	FaceQnet-v1	Sharpness-1	Sharpness-2
Mated-other	PNG-resized	**2**	**2**	**2**	7	8	**2**	13	15
	JPEG	5	4	**0**	5	4	2	69	43
	JPEG 2000	2	**0**	1	1	2	**0**	7	9
	JPEG XL	2	5	2	4	13	**1**	7	14
	Combined	3	3	**1**	4	7	**1**	24	20
Mated-self	PNG-resized	**1**	2	3	7	8	3	13	15
	JPEG	2	**1**	3	3	**1**	5	71	45
	JPEG 2000	**0**	1	2	1	1	1	8	10
	JPEG XL	**1**	4	2	4	14	**1**	9	15
	Combined	**1**	2	3	4	6	2	25	21

Table 1. Distances between normalized quality and mated-other/self comparison score curves from Figure 3, as rounded $[0, 100]$ percentage values (lower is better). For each row the lowest values are printed in **bold**. The "combined" rows show the mean values across the four compression types.

tify how well each FIQA method captures the compression effect on the comparison scores, we computed the distances between these normalized quality/comparison score curves, shown in Table 1. More specifically, we computed the areas under each linearly interpolated curve, and then computed the absolute value of the difference between areas for each pairing of FIQA method and mated-other/mated-self comparison configuration.

Overall, the two sharpness metrics appeared distinctly less reliable than the six modern utility-centric FIQA models. Of those six models, FaceQnet-v0 performed slightly worse overall, and substantially worse for JPEG XL. FaceQnet-v1, which is architecturally identical to FaceQnet-v0, did however perform very well both overall and for JPEG XL, which could indicate that the training procedure might be more important than the architectural differences of these FIQA models, at least in the context of this compression effect scenario. The differences between the results for the remaining FIQA models appear to be comparatively minor overall.

3.3. Additional results

Besides the primary results presented in the previous subsections, we also examined higher target sizes for the ROI images (27kB, 22kB, 17kB, 12kB, 10kB), evaluated the same experiments with a second "portrait" image variant (with target sizes 30kB, 24kB, 18kB, 12kB, 10kB, 9kB, 8kB, 7.7kB), and analyzed a non-mated comparison trial configuration.

For the portrait image variant the ColorFERET images were cropped to the minimal allowed width and height relative to the Inter-Eye Distance (IED) and Eye-Mouth Distance (EMD) as required by ISO/IEC 39794-5:2019 [2]. The cropped portrait region was out of the original image bounds for 43 images, with the out of bounds pixels being filled with black. Other standard requirements were not checked and the IED was below the 90 pixel minimum for 243 images, while an additional 751 images fell below the 120 pixel IED "best practice" value of the standard, so images in this por-

trait variant were only partially compliant. The comparison score results did however not differ substantially between the compression types for the portrait image variant, nor with the higher target sizes for the ROI image variant.

The "non-mated" comparison trial configuration consisted out of randomly selected non-mated [1] trials, equal in number to the mated-other trials (3484), likewise comparing images with the same target size compression applied. Results for this configuration on the ROI images were arguably less interesting for the analysis, due to only comparatively minor changes in the mean comparison scores even at the lower target sizes, with substantial overlap of the computed individual scores and their minima/maxima. We can however note that PNG-resized was comparatively most influential in terms of altering the mean score and the worst-case (maximum) scores, which may indicate that the seemingly competitive mated configuration results for PNG-resized were facilitated in part by removing or "blurring" more recognition-relevant information in a similar manner across images. Nevertheless, the differences were minor, and so PNG-resized appeared to be competitive with JPEG and JPEG 2000 given the results overall, especially at higher target sizes.

4. CONCLUSION

Based on the results we recommend JPEG XL for face image compression across all target sizes, both for face ROI images and portrait images. The differences in the effect on face recognition between compression types did become substantial for face ROI compression at sizes approximately below 5kB. For higher sizes, above 10kB, JPEG XL can still provide the best worst-case results, but the other tested compression types appeared viable at these sizes as well, including the unusual PNG-resized compression. While JPEG compression for ROI images led to slightly better results than JPEG 2000 and PNG-resized for target sizes at and above 3.5kB, it also led to mostly worse results for the very low target sizes 3kB, 2.5kB, and 2.2kB. The advantage of JPEG 2000 over JPEG at lower target sizes supports prior work results mentioned in the introduction [3][6].

Regarding FIQA, the two sharpness measures did not capture the effect of compression on the comparison scores well, especially not for JPEG, while the six modern utility-centric FIQA models mostly did. If only the effect of the recommended JPEG XL compression on the utility were to be measured, then the results indicate FaceQnet-v1, CR-FIQA(S), and MagFace as the primary candidates. To more specifically select one of the six FIQA models for general use, we do however recommend to refer to more general FIQA performance evaluations.

Future work could investigate the qualitative optimization of compression settings, or evaluate the compression speed and assess the potential for computational performance optimizations. Future development work could furthermore attempt to create specialized face image compression types.

5. REFERENCES

[1] ISO/IEC JTC1 SC37 Biometrics, *ISO/IEC 2382-37:2022 Information technology - Vocabulary - Part 37: Biometrics*, International Organization for Standardization, 2022.

[2] ISO/IEC JTC1 SC37 Biometrics, *ISO/IEC 39794-5:2019 Information technology - Extensible biometric data interchange formats - Part 5: Face image data*, International Organization for Standardization, 2019.

[3] W. Funk, M. Arnold, C. Busch, and A. Munde, "Evaluation of Image Compression Algorithms for Fingerprint and Face Recognition," in *Proc. IEEE Information Assurance Workshop*, 2005.

[4] J. Phillips, H. Wechsler, J. Huang, and P. Rauss, "The FERET Database and Evaluation Procedure for Face Recognition Algorithms," *Image and Vision Computing Journal*, vol. 16, no. 5, pp. 295–306, 1998.

[5] J. Phillips, H. Moon, S. Rizvi, and P. Rauss, "The FERET Evaluation Methodology for Face-Recognition Algorithms," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, pp. 1090–1104, October 2000.

[6] K. Delac, M. Grgic, and S. Grgic, "Image Compression Effects in Face Recognition Systems," in *Face Recognition*, chapter 5. IntechOpen, 2007.

[7] J. Deng, J. Guo, and S. Zafeiriou, "ArcFace: Additive Angular Margin Loss for Deep Face Recognition," in *Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[8] ISO/IEC JTC1 SC29 Coding of audio, picture, multimedia and hypermedia information, *ISO/IEC 18181-1:2022 Information technology - JPEG XL image coding system - Part 1: Core coding system*, International Organization for Standardization, 2022.

[9] ISO/IEC JTC1 SC29 Coding of audio, picture, multimedia and hypermedia information, *ISO/IEC 10918-1:1994 Information technology - Digital compression and coding of continuous-tone still images: Requirements and guidelines*, International Organization for Standardization, 1994.

[10] ISO/IEC JTC1 SC29 Coding of audio, picture, multimedia and hypermedia information, *ISO/IEC 15444-1:2019 Information technology - JPEG 2000 image coding system - Part 1: Core coding system*, International Organization for Standardization, 2019.

[11] ISO/IEC JTC1 SC24 Computer graphics, image processing and environmental data representation, *ISO/IEC 15948:2004 Information technology - Computer graphics and image processing - Portable Network Graphics (PNG): Functional specification*, International Organization for Standardization, 2004.

[12] T. Schlett, C. Rathgeb, O. Henniger, J. Galbally, J. Fierrez, and C. Busch, "Face Image Quality Assessment: A Literature Survey," *ACM Computing Surveys (CSUR)*, December 2021.

[13] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments," Tech. Rep., University of Massachusetts, Amherst, October 2007.

[14] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou, "RetinaFace: Single-Shot Multi-Level Face Localisation in the Wild," in *Conf. on Computer Vision and Pattern Recognition (CVPR)*. June 2020, pp. 5202–5211, IEEE.

[15] F. Boutros, M. Fang, M. Klemt, B. Fu, and N.Damer, "CR-FIQA: Face Image Quality Assessment by Learning Sample Relative Classifiability," 2021.

[16] D. Yi, Z. Lei, S. Liao, and S. Z. Li, "Learning Face Representation from Scratch," November 2014.

[17] Q. Meng, S. Zhao, Z. Huang, and F. Zhou, "MagFace: A Universal Representation for Face Recognition and Quality Assessment," in *Conf. on Computer Vision and Pattern Recognition (CVPR)*. June 2021, pp. 14225–14234, IEEE.

[18] P. Terhörst, J. N. Kolf, N. Damer, F. Kirchbuchner, and A. Kuijper, "SER-FIQ: Unsupervised Estimation of Face Image Quality Based on Stochastic Embedding Robustness," in *Conf. on Computer Vision and Pattern Recognition (CVPR)*. June 2020, pp. 5650–5659, IEEE.

[19] J. Hernandez-Ortega, J. Galbally, J. Fierrez, R. Haraksim, and L. Beslay, "FaceQnet: Quality Assessment for Face Recognition based on Deep Learning," April 2019.

[20] J. Hernandez-Ortega, J. Galbally, J. Fierrez, and L. Beslay, "Biometric Quality: Review and Application to Face Recognition with FaceQnet," June 2020.

[21] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VGGFace2: A Dataset for Recognising Faces across Pose and Age," *Intl. Conf. on Automatic Face and Gesture Recognition*, May 2018.

[22] F. Crété, T. Dolmiere, P. Ladret, and M. Nicolas, "The Blur Effect: Perception and Estimation with a New No-Reference Perceptual Blur Metric," in *SPIE Electronic Imaging Symposium Conf. Human Vision and Electronic Imaging*, San Jose, United States, Jan 2007, vol. XII.