

Identity Preserving Loss for Learned Image Compression

Jiuhong Xiao ^{*1,2}, Lavisha Aggarwal², Prithviraj Banerjee², Manoj Aggarwal², and Gerard Medioni²

¹New York University ²Amazon

jx1190@nyu.edu aggarl@amazon.com banerjee@amazon.com manojagg@amazon.com
medioni@amazon.com

Abstract

Deep learning model inference on embedded devices is challenging due to the limited availability of computation resources. A popular alternative is to perform model inference on the cloud, which requires transmitting images from the embedded device to the cloud. Image compression techniques are commonly employed in such cloud-based architectures to reduce transmission latency over low bandwidth networks. This work proposes an end-to-end image compression framework that learns domain-specific features to achieve higher compression ratios than standard HEVC/JPEG compression techniques while maintaining accuracy on downstream tasks (e.g., recognition). Our framework does not require fine-tuning of the downstream task, which allows us to drop-in any off-the-shelf downstream task model without retraining. We choose faces as an application domain due to the ready availability of datasets and off-the-shelf recognition models as representative downstream tasks. We present a novel Identity Preserving Reconstruction (IPR) loss function which achieves Bits-Per-Pixel (BPP) values that are $\sim 38\%$ and $\sim 42\%$ of CRF-23 HEVC compression for LFW (low-resolution) and CelebA-HQ (high-resolution) datasets, respectively, while maintaining parity in recognition accuracy. The superior compression ratio is achieved as the model learns to retain the domain-specific features (e.g., facial features) while sacrificing details in the background. Furthermore, images reconstructed by our proposed compression model are robust to changes in downstream model architectures. We show at-par recognition performance on the LFW dataset with an unseen recognition model while retaining a lower BPP value of $\sim 38\%$ of CRF-23 HEVC compression.

1. Introduction

There is an increasing trend to perform tasks like object detection and semantic segmentation on mobile or resource-constrained local devices. However, at-edge computing often cannot afford the high computation complexity of large CNN or Transformer models. Instead of reducing model

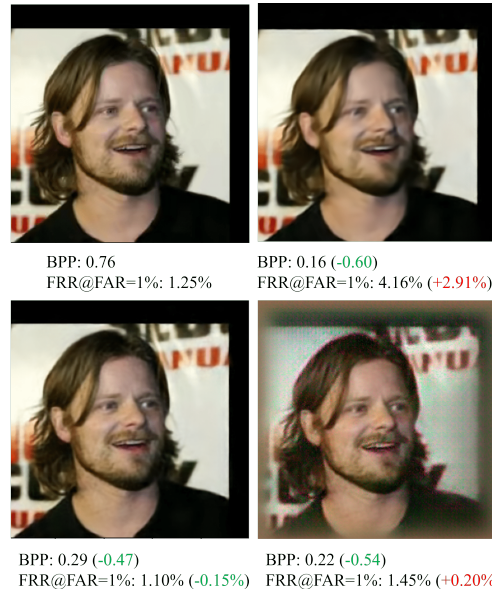


Figure 1. Identity preserving loss for learned image compression. Top left: HEVC baseline; Top right: The result shows that if the compression model only optimizes rate-distortion trade-off, the significant distortion appears on the face region and degrades recognition performance; Bottom left: The compression model optimizing Identity Preserving Reconstruction (IPR) loss can maintain image quality with improved recognition accuracy; Bottom right: If the compression model optimizes Identity Preserving (IP) loss, it sacrifices the image quality of background and maintains the clear face region with slightly lower accuracy and much lower BPP than baseline.

complexity on local devices, a popular solution is to transfer images from the device to a remote server and execute computational-intensive tasks on the cloud.

This paper focuses on an end-to-end image compression framework for the face recognition domain due to the ready availability of datasets, and off-the-shelf recognition models as representative downstream tasks. For face-related tasks, recent works focus on high-resolution face images for high-precision face recognition [16] and 3D face reconstruction [9]. Local devices with limited resources will show high inference latency on these tasks. There is also a potential requirement to store raw images on database

*Work done during an internship at Amazon.

servers for the downstream pipeline. Under these circumstances, a remote server needs to receive and store high-resolution images. Image compression methods help reduce the transmission latency and storage space on the server.

Standard lossy compression techniques like JPEG [29] and HEVC [24] reduce image file size without noticeable distortion on images. With the growing popularity of deep learning [17], an increasing number of recently published works are using CNN models to compress images to achieve a higher compression ratio and less notable compression artifacts [1, 2, 21, 26]. Further works [5, 6] indicate that end-to-end image compression training with downstream tasks (like object detection) can highlight regions relevant to the downstream task, and further increase the compression ratio with improved downstream performance.

Since our compression model only receives face images, it allows us to tune the compression model against the face domain and further increase the compression ratio. We regard our compression model as a plug-in module between the image capture module and the downstream system, so we update only the compression model during training. We use HEVC(CRF=23) and JPEG as standard compression baselines, and provide a performance comparison for a) low-resolution vs. high-resolution, and b) unaligned vs. aligned face datasets.

Our method has multiple advantages: First, our method does not require additional retraining or fine-tuning of the downstream recognition model. Second, we show that the compression models trained for a specific domain (faces in our case) show a higher compression ratio and recognition accuracy than those trained with generic images. Third, a new framework with Identity Preserving Reconstruction (IPR) loss or Identity Preserving (IP) loss is proposed to improve the recognition performance of the images generated from the compression model. The results show that all models with IPR loss have improved recognition performance for the low-resolution dataset. Our method achieves a high compression ratio with slightly lower recognition accuracy for the high-resolution dataset. Lastly, we present an evaluation of our proposed algorithm’s robustness to downstream model changes, and show improved recognition accuracy and high compression ratio with an unseen recognition model.

2. Related Work

For lossy compression, a learnable image compression method should consider both the network backbone and the entropy model [11]. There are two approaches for the network backbone: autoencoder-like architectures [2, 14, 21] and recurrent frameworks [26]. Autoencoder-like architectures can reduce dimensionality to produce a compact latent representation. Recurrent frameworks are used to encode the residue to compress the image progressively.

For the entropy model, there are several options like context-based prediction [2, 22] or calculating histogram of latent representation [25]. One of the most widely-used frameworks, the hyper-prior method [2], follows the VAE architecture and uses the hyper-prior model as an entropy model. From the survey [11], the authors observe that the hyper-prior [2] method outperforms other methods for its outstanding entropy model.

There are multiple variants of the hyper-prior method: [14] conducts network architecture search (NAS) on the hyper-prior model to reduce the amount of computation, while not reducing the encoded file size. Alternately, [21] integrates GAN with the hyper-prior model to generate images that achieve a high subjective image quality (realistic texture). However, such a method encourages “photo-realism” over preserving the fine information detail in the images, which may impact downstream tasks that need original details (such as recognition algorithms).

For task-specific image compression, the work [5, 6] shows that by jointly optimizing the compression model and an object detection model, the compressed images highlight the task-related region and do not decrease the downstream performance significantly. Some works [7, 19] also perform task-based JPEG compression. Recent work [28] points out that task-specific compression can be done without decoding for classification and semantic segmentation tasks by jointly training both the decoder and the classifier using intermediate code. Compared to these works, our method applies the task-specific image compression method to the face imagery domain, without any requirements on optimizing/fine-tuning the downstream recognition model, and the encoder-decoder-recognition paradigm instead of the encoder-recognition [28] paradigm for increased generalizability to different downstream recognition models.

For face-specific learnable compression, [31] uses deep-feature-based compression. It applies compression to the feature map extracted by FaceNet [23]. Deep-feature-based compression only allows compression for aligned images and depends on the FaceNet feature tuned to face recognition. [32] uses a talking-head video synthesis model to allow face-specific video compression. Our method focuses on both unaligned and aligned still images and can be generalized to other tasks than recognition. Besides, [31] uses an unlearned entropy model, while our method uses a learned entropy model, which allows a higher compression ratio. Overall, our method can show higher flexibility, and compression performance than [31].

3. Compression for Face Recognition

Our proposed framework is shown in Fig. 2. Images x are first encoded and quantized to discrete latent representation \hat{y} . We follow the hyper-prior method [2] to set the entropy model. The hyper encoder-decoder outputs discrete

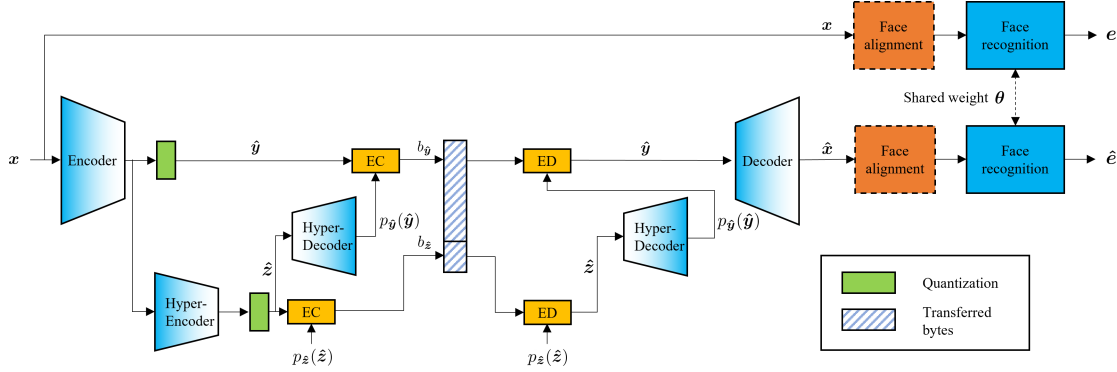


Figure 2. Framework for the compression model with face recognition system. The face alignment is optional depending on if the images are pre-aligned.

hyper-latent representation \hat{z} and predicts the distribution of \hat{y} using the entropy model $p_{\hat{y}}(\hat{y})$. An entropy encoder (EC) and an entropy decoder (ED) are for lossless compression. For EC, $p_{\hat{y}}(\hat{y})$ is used to encode \hat{y} to transferable bytes $b_{\hat{y}}$ and a separate entropy model $p_{\hat{z}}(\hat{z})$ with a factorized prior distribution is used to encode \hat{z} to $b_{\hat{z}}$. $b_{\hat{y}}$ and $b_{\hat{z}}$ make up the payload that needs transmission over the network. After transmission, ED will transform those bytes back to discrete latent and hyper-latent representation. Finally, the decoder reconstructs the images \hat{x} .

The existing compression methods optimize reconstruction loss to maintain image quality and bit-rate loss to reduce the file size:

$$\mathcal{L}_R = \mathcal{L}_{\text{rec}} + \lambda_{\text{rate}} \mathcal{L}_{\text{rate}} \quad (1)$$

$$\mathcal{L}_{\text{rec}} = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} [d(\mathbf{x}, \hat{\mathbf{x}})] \quad (2)$$

$$\mathcal{L}_{\text{rate}} = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} [-\log(p_{\hat{y}}(\hat{y})) - \log(p_{\hat{z}}(\hat{z}))] \quad (3)$$

where \mathcal{L}_{rec} is the reconstruction loss, $\mathcal{L}_{\text{rate}}$ is the bit-rate loss, $d(\mathbf{x}, \hat{\mathbf{x}})$ is the distortion function, and λ_{rate} is the weight for bit-rate loss. We name the total loss function \mathcal{L}_R as **Reconstruction-only** loss. When selecting the MSE function as the reconstruction loss, $d(\mathbf{x}, \hat{\mathbf{x}})$ is defined as:

$$d(\mathbf{x}, \hat{\mathbf{x}}) = \|\mathbf{x} - \hat{\mathbf{x}}\|^2 \quad (4)$$

Other options are the MS-SSIM function [33] or some perceptual metrics that can measure the distortion between original images \mathbf{x} and reconstructed images $\hat{\mathbf{x}}$.

To make quantization differentiable in training, the hyper-prior method [2] adds uniform noise to simulate quantization error and calculate $\mathcal{L}_{\text{rate}}$. However, in [25], using additive noise as the quantization alternative adds strong visual artifacts into decompressed images for JPEG compression. Therefore, we still use (hyper-)latent variables (\hat{y} and \hat{z}) with additive noise when calculating $p_{\hat{y}}(\hat{y})$, $p_{\hat{z}}(\hat{z})$

for $\mathcal{L}_{\text{rate}}$, but send the (hyper-)latent variables quantized by rounding function (used in the actual encoding process) to (hyper-)decoder, using a differentiable back-propagation function like [21].

3.1. Identity Preserving Reconstruction Loss

Our goal is to improve the compression ratio without degrading recognition performance. For this task, we use the CurricularFace [13] pre-trained model. We do not fine-tune the model weights (we freeze the weights instead).

To make the compression model compatible with the recognition model, we derive the **Identity Preserving (IP)** loss function, including the bit-rate loss and identity loss to reduce the distance between e and \hat{e} (Fig. 2) for identity recognition with smaller file size:

$$\mathcal{L}_{\text{IP}} = \lambda_{\text{rate}} \mathcal{L}_{\text{rate}} + \lambda_{\text{id}} \mathcal{L}_{\text{id}} \quad (5)$$

$$\mathcal{L}_{\text{id}} = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} \left[1 - \frac{e \cdot \hat{e}}{\|e\| \|\hat{e}\|} \right] \quad (6)$$

where \mathcal{L}_{id} is the identity loss. Since \mathcal{L}_{IP} only optimizes recognition performance, we further derive **Identity Preserving Reconstruction (IPR)** loss function to jointly optimize image quality and recognition performance:

$$\mathcal{L}_{\text{IPR}} = \mathcal{L}_{\text{rec}} + \lambda_{\text{rate}} \mathcal{L}_{\text{rate}} + \lambda_{\text{id}} \mathcal{L}_{\text{id}} \quad (7)$$

We infer the embeddings from decompressed images and original images via face alignment module A and face recognition model R with shared weights θ as

$$e = R_{\theta}(A(\mathbf{x})) \quad \text{and} \quad \hat{e} = R_{\theta}(A(\hat{\mathbf{x}})) \quad (8)$$

Since we use the embeddings to match images with identities, if the cosine distance between \hat{e} and e is small enough, decompressed images are likely to maintain parity in recognition performance.

To simplify the hyperparameter fine-tuning, we set $\lambda_{id}=1.0$ or 0.0 and only fine-tuned λ_{rate} . IPR loss with $\lambda_{id}=0.0$ is equivalent to Reconstruction-only loss. Simultaneously fine-tuning λ_{rate} and λ_{id} may help the model get a better balance between multiple objectives, but that is not the focus of this work.

Distinguished from other task-specific or face-specific compression methods, our method considers to:

- **Compress unaligned images** Although the recognition performance is only related to the aligned face region, our method compresses either unaligned or aligned images. It uses reconstruction loss to optimize the overall image quality, not only for face regions. Compressing unaligned images is for practical retention decisions to save unaligned images for reference, and allowance for changing the alignment methods.
- **Avoid fine-tuning downstream models** Our method does not (end-to-end) fine-tune or retrain the downstream model by any face recognition loss. First, this can avoid additional effort to fine-tune hyperparameters of the recognition loss and reduce the data scale and training time required for fine-tuning a high-complexity face recognition model. Second, this can decouple the compression and downstream models. We can use the same downstream model to achieve variable bit rates (training compression models with different hyperparameters). Third, this helps maintain the robustness to downstream model change. When end-to-end fine-tuning the recognition model, the recognition model will adapt to the compressed images, allowing more unexpected distortion from the compression model. The exceeded distortion may degrade the recognition performance when another recognition model (trained on raw face images) is attached but not retrained on compressed images.

3.2. Face Alignment Module

For the framework (Fig. 2), we attach the face alignment module and face recognition model to calculate the embeddings. For alignment, we use the MTCNN model from FaceNet [23] to detect faces. It outputs the bounding boxes and 5-point landmarks for each face. We select the face that has the largest size and locates centrally.

To compare different alignment orders in chosen datasets, we will use CelebA-HQ and CelebA-HQ-align mentioned in Section 4.1. We do post-alignment and pre-alignment on these datasets, respectively.

Since our pre-trained model only accepts 112x112 images, we need to resize images for CelebA-HQ related high-resolution images after the alignment. Fig. 3 shows the additional resizing operation on aligned high-resolution face




Figure 3. The results of alignment and resizing operation (from left to right). LFW images are aligned from 250x250 to 112x112. CelebA-HQ images are aligned from 1024x1024 to 448x448 and then further resized to 112x112.

images, while the low-resolution face images only require an alignment for recognition.

4. Results

4.1. Datasets

Table 1 shows the datasets for evaluating different image compression approaches. Images for all chosen datasets are available only in JPEG format; hence we use those as original images without compression. We use the LFW [12] dataset as a low-resolution dataset, and CelebA-HQ [15] dataset as a high-resolution dataset (We only use a subset of faces that are also used by CelebAMask-HQ [18]). Since CelebA-HQ dataset does not have an official train/test split, we create a split to make the number of train identities approximately equal to the number of test identities. We use the face alignment algorithm mentioned in Section 3.2 to align the images of CelebA-HQ and create CelebA-HQ-align. CelebA-HQ-align uses pre-alignment, and the image size is 448x448. We also use CLIC dataset [27] to train a general compression model for comparison.

4.2. Metrics

To evaluate the performance of our image compression models, we will be utilizing the following metrics.

4.2.1 Evaluating Compression Ratio

We use the bits per pixel (BPP) to measure the performance of the compression ratio. The BPP is measured as

$$\text{BPP} = \frac{\text{File size in bits}}{\# \text{ pixels in one image}} \quad (9)$$

We target low BPPs in this work (2x to 5x compression ratio of CRF-23-HEVC files).

Dataset Name	#Train Images	#Train Identities	#Test Images	#Test Identities	Raw image size
LFW	9k	4k	4k	2k	250x250
CelebA-HQ	20k	3k	10k	3k	1024x1024
CelebA-HQ-align	20k	3k	10k	3k	448x448
CLIC	626	N/A	60	N/A	Unspecified

Table 1. Dataset Information. Training on CLIC images uses small patches from the original images since the original images are too large. Training on other datasets is based on the whole image.

4.2.2 Evaluating Recognition Performance

For evaluating the accuracy of recognition tasks, we use $FRR@FAR=1\%$ metric. The metrics are measured using the following process:

1. Put N images per identity into the gallery set and the rest into the query set.
2. For each image in the query set, measure the embedding distance between this image and all images in the gallery set.
3. Label the lowest embedding distance from the same identity as *same distance*, and the lowest embedding distance from any other not-same identities as *not-same distance* for identity matching.
4. Plot the ROC curve for all images in the query set and calculate $FRR@FAR=1\%$ (false rejection rate when false acceptance rate equals to 1%).

For the LFW dataset, we set $N = 1$ to evaluate the recognition performance. Our recognition model is trained to have high accuracy on LFW and is not fine-tuned on CelebA-HQ. Hence for CelebA-HQ and CelebA-HQ-align datasets, we set $N = 3$ to have improved representation for each identity in the gallery set. Note that our evaluation protocol is more challenging compared to the standard LFW face verification protocol [12]. Our protocol uses a more comprehensive negative pair set to compare each query against the entire gallery. In contrast, the standard protocol uses a random sampling process to generate a limited set of image pairs. Our primary goal is to maintain recognition parity with CRF-23-HEVC compressed images while achieving higher compression ratios.

4.2.3 Evaluating Image Quality

We evaluate image quality in terms of Peak signal-to-noise ratio (PSNR) and Multi-scale structural similarity index measure (MS-SSIM) [33] metrics. The goal is to achieve a similar image quality to HEVC compressed images w.r.t. PSNR and MS-SSIM metrics when using Reconstruction-only loss or IPR loss (IP loss does not optimize image quality).

4.3. Experiment Settings

To compare with standard compression techniques, we use HEVC, one of the latest state-of-the-art coding methods, with CRF=23 as our main baseline. We are focusing on HEVC intraframe compression performance on still images. The video compression performance is not evaluated since the embedding system for face recognition rarely transfers the whole video for latency consideration. The compression quality is set to balance the BPP and recognition performance trade-off. We also present the results for JPEG series. Since the original format is JPEG, we recompress the images at varying JPEG compression levels.

We trained models with different combinations of reconstruction and identity preserving loss functions: Reconstruction-only loss (equation 1, **L2_REC** and **MS-SSIM_REC**), IPR loss (equation 7, **L2_IPR** and **MS-SSIM_IPR**), IP loss (equation 5, **L2_IP** and **MS-SSIM_IP**). The prefixes “L2” or “MS-SSIM” indicate the form of reconstruction loss functions (equation 2) we use for training. For the models trained with IPR or IP loss function, the weights are initialized from those trained with Reconstruction-only loss.

We also train a general image compression model, where we use the same structure as our framework without face alignment and face recognition modules. The training images are the small patches from the original images (Table 1), and we apply L2 loss as the reconstruction loss function. We name the model as **L2_baseline**.

When training with a recognition model, we use a fixed CurricularFace [13] face recognition model (pre-trained on refined MS1M-v2 dataset [8] and SOTA on IJB-C dataset [20]). During the evaluation, we use the CurricularFace model to compare the performance and a CosFace [30] (pre-trained on VGGFace2 dataset [4]) model to evaluate the robustness of our framework on an unseen recognition model.

4.4. Quantitative Results

In the following sections, we discuss our quantitative results on LFW dataset (Fig. 4) and CelebA-HQ(-align) datasets (Fig. 5). We have the following key findings in our experiments.




Figure 4. The image quality and recognition performance of JPEG, HEVC, and our method on the LFW dataset. From left to right: Image quality results (PSNR and MS-SSIM), recognition results with CurricularFace model, recognition results with CosFace model.

4.4.1 IPR Loss Balances the Trade-off between Image Quality, Recognition Performance, and File Size

In Fig. 4, We show image quality performance and recognition performance on models trained with Reconstruction-only Loss, IPR loss, and IP loss on the LFW dataset.

As the results show, models trained with Reconstruction-only loss (L2_REC and MS-SSIM_REC) mostly have better image quality performance than those trained with IPR or IP loss. For example, L2_REC has 1.43 more PSNR than L2_IPR, and 18.45 more PSNR than L2_IP when BPP is about 0.3. However, L2_REC and MS-SSIM_REC have worse recognition performance (FRR@FAR=1% increases to 4.16% for L2_REC) when the BPPs are less than 0.25.

Models trained with IPR loss (L2_IPR and MS-SSIM_IPR) improve recognition performance without significantly decreasing image quality. For instance, L2_IPR has only 1.43 less PSNR but 0.4% less FRR@FAR=1% than L2_REC when BPP is about 0.3. But training with IPR loss leads to slightly higher BPP. It moves curves to the right side in Fig. 4, and the lowest BPP is 0.29.

Models trained with IP loss show the path to lower BPPs with acceptable recognition performance. MS-SSIM_IP (BPP=0.2) increases only 0.05% FRR@FAR=1% than HEVC, but gets 73% less BPP than HEVC. However, these models show a significant deterioration in image quality metrics. When BPP is about 0.3, L2_IP has 18.45 less PSNR and 0.097 less MS-SSIM than L2_REC, and MS-SSIM_IP has 14.34 less PSNR and 0.102 less MS-SSIM than MS-SSIM_REC.

We also investigate the trade-off between image quality performance, recognition performance, and BPP in Fig. 4: Models trained on IP loss or Reconstruction-only loss can reach low BPPs, which are less than 0.25. They can also get a better recognition performance or image quality than each other, respectively. L2_IP has 0.25% less FRR@FAR=1% and 18.45 less PSNR than L2_REC when BPP is about 0.3. When models are trained with IPR loss (L2_IPR and MS-SSIM_IPR), image quality and recognition performance are improved, but the resulting BPPs are larger. We find no

models with IPR loss can reach BPP=0.25. We also observe that when the BPPs decrease below 0.25, the models with IP loss (purple and gray curves) have a sharp deterioration in FRR@FAR=1% values. This may indicate a lower bound of BPP to maintain parity on recognition performance.

We summarize that one can focus on either image quality (L2_REC and MS-SSIM_REC) or recognition performance (L2_IP and MS-SSIM_IP), but not both when trying to have a BPP lower than 0.25. When the BPP is lower than 0.25 by focusing on recognition performance, models with IP loss will start to degrade in recognition performance.

Discussion on performance trends In Fig. 5, we observe that for high-resolution datasets (CelebA-HQ and CelebA-HQ-align), the performance trend is unstable. Lower BPP can get better performance. The advantage of IPR loss is less significant than that of the LFW dataset. This may be due to the poor generalization of the pre-trained recognition model on these datasets. The distribution histogram of the recognition embedding distance shows the recognition performance on original images is poor: There are a lot of same-identity pairs with high distance values, indicating that the recognition model cannot correctly match these images with the identities. This error could introduce noise as the IPR loss uses this feedback from the recognition model to generate decompressed images.

Another possible reason might be the resizing operation in the face alignment module on the high-resolution datasets. The face alignment module for CelebA-HQ(-aligned) images resizes aligned images from (448, 448) to (112, 112) to fit the recognition model’s input resolution (as shown in Fig. 5). In this case, IPR loss is more likely to ignore artifacts on the face region (also discussed in Section 4.5.2).

In Fig. 4, for the low-resolution dataset (LFW), thanks to avoiding both reasons listed above, the noise becomes much smaller but still exists (for L2_IPR and MS-SSIM_IPR curves). This may be caused by the noise from balancing the tradeoff between image quality and recognition performance.




Figure 5. Image quality and recognition performance of JPEG, HEVC and our methods on CelebA-HQ (left two plots) and CelebA-HQ-align (right two plots) with CurricularFace model

4.4.2 Training against Face Domain Enables Higher Compression Ratio

To compare the performance between general compression models and domain-specific compression models, we trained the general compression model on the CLIC training set. We evaluated it on the CelebA-HQ test set (Fig. 5). Both datasets have similar resolutions to reduce the influence of scale change. The results show that our models trained on the CelebA-HQ dataset (L2_REC) have better image quality performance and recognition accuracy. For example, when BPP is about 0.06, L2_REC has 3.08 more PSNR and 0.67% less FRR@FAR=1% than L2_baseline. Comparing models with about 35.50 PSNR, L2_REC has 29% less BPP and 1.15% less FRR@FAR=1% than L2_baseline. This shows that when we focus on the face domain, compression models trained against face images benefit the compression ratio without decreasing image quality and recognition accuracy.

We also find a large gap between expected BPPs and actual BPPs of the general compression model on face datasets. In Fig. 5, for the leftmost point of the L2_baseline curve, when we set the target BPP to be 0.100, the validation BPP on the CLIC test set is 0.095, which is close to the target, but the test BPP on CelebA-HQ test set is 0.069, which shows a non-trivial decrease on the BPP. For this unexpected low BPP, the compressed images have a low PSNR score of 31.89. The drop of BPPs may be caused by the domain gap between CLIC and CelebA-HQ. It increases the uncertainty in compression ratio and image quality if the general model is applied to a new dataset on specific domains and will be an unstable factor in transmission.

4.4.3 Robustness to Downstream Model Change

In Fig. 4, we show the recognition performance of our method when training on CurricularFace but evaluating on another recognition model CosFace [30]. In the results, MS-SSIM_IP (BPP=0.3) shows 0.11% less FRR@FAR=1% and 60% less BPP than HEVC result. L2_IPR (BPP=0.29)

shows 0.76% less FRR@FAR=1% and 62% less BPP than HEVC result. The results with IPR loss or IP loss show that the compression model trained with one recognition model can have good generalization performance with another recognition model on the LFW dataset.

We have shown that models with IP loss (L2_IP and MS-SSIM_IP) have significant distortion (shown in decreased PSNR in Fig. 4) on the images when the quality level is set to the lowest (leftmost points of the curves). But still, the models perform at most 0.35% more FRR@FAR=1% than HEVC baseline with the CurricularFace model. This implies that IP loss corrupts only task-unrelated regions but maintains recognizable face features. The result with the CosFace model also depicts that good generalization performance may benefit from these maintained face features.

4.5. Qualitative Results

In this section, we present qualitative visualizations imagery from LFW and CelebA-HQ(-align) datasets. Fig. 6-7 show the reconstructed images with typical compression methods and qualities. All the images use the lowest quality level (the leftmost points of curves) to better visualize possible artifacts. We make the following observations:

4.5.1 IPR loss vs. IP loss

For models with IPR loss, the results of L2_IPR and MS-SSIM_IPR do not show notable artifacts on the images. However, models with IP loss (L2_IP and MS-SSIM_IP) have strong distortion on the images. The difference between IPR loss and IP loss is the reconstruction loss term \mathcal{L}_{rec} . Models with IPR loss optimize reconstruction loss to maintain high image quality, but models with IP loss allow more distortion if the predicted embedding \hat{e} is close enough to the original embedding e .

Discussion on artifacts from IP loss Models with IP loss show different artifacts on unaligned datasets (LFW and CelebA-HQ) and the aligned dataset (CelebA-HQ-align). The unaligned compressed images show the corrupted boundary and distorted background. At the same

time, in aligned datasets (CelebA-HQ-align), the visualized results have mesh artifacts on the whole image, where the face region is strongly distorted. The influence of the face alignment module on unaligned images is shown: Face alignment will discard the boundary and only crop the center face region of the images. Therefore, it allows the compression model with IP loss to focus on the center face region and reduces bitrate for the boundary and background.

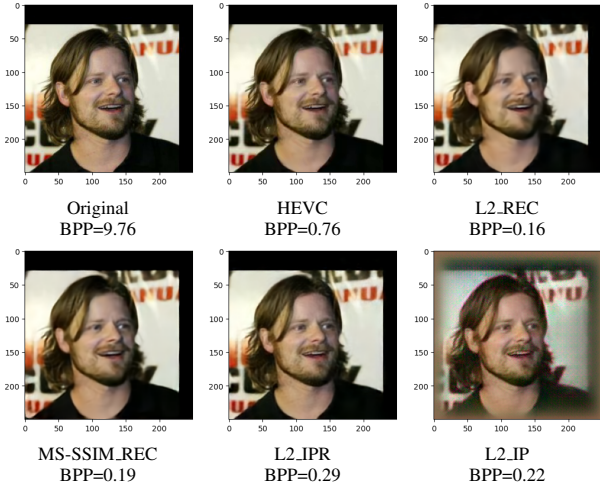


Figure 6. Visualization results on LFW datasets




Figure 7. Visualization results on CelebA-HQ (Upper) and CelebA-HQ-align (Lower) datasets

Discussion with mask-based compression The result of IP loss is visually similar to mask-based compression methods [3, 10] on LFW and CelebA-HQ datasets, which explicitly or implicitly mask out the background region and allocate more bits on the ROI to maintain its parity. These methods require explicit ROI annotations, which have more annotation costs than our method. There can be a mismatch

between ROIs expected by humans and the feature regions that downstream models are using. Our method avoids this mismatch by directly using the feedback from downstream models to adjust the allocation of bits.

4.5.2 High-resolution vs. Low-resolution

In the LFW low-resolution dataset, L2_IPR and L2_IP are shown to produce sharp and clear face regions. In CelebA-HQ and CelebA-HQ-align high-resolution datasets, all models have blurring artifacts and striped texture on the face region when the BPPs are low than 0.075 (CelebA-HQ) and 0.15 (CelebA-HQ-align). Even though the face region is not so clear, the recognition accuracy is maintained: L2_REC has 0.46% less FRR@FAR=1% for CelebA-HQ and 1.15% less FRR@FAR=1% for CelebA-HQ-align than HEVC baseline.

The influence of the resizing operation on high-resolution datasets has been discussed in 4.4.1, and the visualization results support the point. The pre-trained recognition model requires CelebA-HQ images to be resized, as shown in Fig. 3. Therefore, some artifacts like blurring or striped texture will be invisible to the recognition model. After resizing, the vanishing of artifacts can make the IPR or IP loss ignore these artifacts.

Overall, we show that resizing mitigates the influence of artifacts on recognition performance. Therefore, IPR loss and IP loss are less likely to maintain the details in the face region on high-resolution datasets.

5. Conclusion

In this work, we presented an end-to-end image compression framework trained to leverage domain-specific features to achieve higher compression ratios while maintaining accuracy on downstream tasks. We presented an Identity Preserving Reconstruction (IPR) loss for the face recognition domain, which achieves BPP values $\sim 38\%$ and $\sim 42\%$ of CRF-23 HEVC compression for LFW and CelebA-HQ datasets, respectively, while maintaining parity in recognition accuracy. Our approach does not require fine-tuning the downstream task model (face recognition model). We also evaluated our approach to unseen downstream recognition models, highlighting the generalizability and robustness of the compressed imagery generated by our proposed framework. In conclusion, our work shows that the proposed framework with IPR loss can guide compression models to improve compression ratios by focusing on essential image features for downstream task accuracy.

References

- [1] E. Agustsson, Fabian Mentzer, M. Tschannen, L. Cavigelli, R. Timofte, L. Benini, and L. Gool. Soft-to-hard vector

- quantization for end-to-end learning compressible representations. In *NIPS*, 2017. 2
- [2] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *International Conference on Learning Representations*, 2018. 2, 3
- [3] Chunlei Cai, Li Chen, Xiaoyun Zhang, and Zhiyong Gao. End-to-end optimized roi image compression. *IEEE Transactions on Image Processing*, 29:3442–3457, 2020. 8
- [4] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 67–74, Los Alamitos, CA, USA, may 2018. IEEE Computer Society. 5
- [5] Lahiru D. Chamain, Fabien Racapé, Jean Begaint, Akshay Pushparaja, and Simon Feltman. End-to-end optimized image compression for machines, a study. *2021 Data Compression Conference (DCC)*, pages 163–172, 2021. 2
- [6] Lahiru D Chamain, Fabien Racapé, Jean Bégaint, Akshay Pushparaja, and Simon Feltman. End-to-end optimized image compression for multiple machine tasks. *arXiv preprint arXiv:2103.04178*, 2021. 2
- [7] Jinyoung Choi and Bohyung Han. Task-aware quantization network for jpeg image compression. In *European Conference on Computer Vision*, pages 309–324. Springer, 2020. 2
- [8] 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. 5
- [9] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Trans. Graph.*, 40(4), July 2021. 1
- [10] S.B. Gokturk, C. Tomasi, B. Girod, and C. Beaulieu. Medical image compression based on region of interest, with application to colon ct images. In *2001 Conference Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 3, pages 2453–2456 vol.3, 2001. 8
- [11] Y. Hu, W. Yang, Z. Ma, and J. Liu. Learning end-to-end lossy image compression: A benchmark. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (01):1–1, mar 5555. 2
- [12] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07–49, University of Massachusetts, Amherst, October 2007. 4, 5
- [13] Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, Jilin Li, and Feiyue Huang. Curricularface: adaptive curriculum learning loss for deep face recognition. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5901–5910, 2020. 3, 5
- [14] Nick Johnston, Elad Eban, Ariel Gordon, and Johannes Ballé. Computationally efficient neural image compression. *arXiv preprint arXiv:1912.08771*, 2019. 2
- [15] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. 4
- [16] Shun-Cheung Lai, Minna Kong, Kin-Man Lam, and Dong Li. High-resolution face recognition via deep pore-feature matching. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 3477–3481, 2019. 1
- [17] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015. 2
- [18] Cheng-Han Lee, Ziwei Liu, Lingyu Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5549–5558, 2020. 4
- [19] Yuzhang Lin, Ashok Amit, Michael Marcellin, and Ali Bilgin. Task-based jpeg 2000 image compression: An information-theoretic approach. In *2018 Data Compression Conference*, pages 423–423, 2018. 2
- [20] Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K Jain, W Tyler Niggel, Janet Anderson, Jordan Cheney, et al. Iarpa janus benchmark-c: Face dataset and protocol. In *2018 International Conference on Biometrics (ICB)*, pages 158–165. IEEE, 2018. 5
- [21] Fabian Mentzer, George D Toderici, Michael Tschannen, and Eirikur Agustsson. High-fidelity generative image compression. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11913–11924. Curran Associates, Inc., 2020. 2, 3
- [22] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. 2
- [23] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015. 2, 4
- [24] Gary J. Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (hevc) standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(12):1649–1668, 2012. 2
- [25] Lucas Theis, W. Shi, A. Cunningham, and Ferenc Huszár. Lossy image compression with compressive autoencoders. *ArXiv*, abs/1703.00395, 2017. 2, 3
- [26] George Toderici, Sean M. O’Malley, Sung Jin Hwang, Damien Vincent, David Minnen, Shumeet Baluja, Michele Covell, and Rahul Sukthankar. Variable rate image compression with recurrent neural networks. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto*

Rico, May 2-4, 2016, Conference Track Proceedings, 2016. [2](#)

- [27] George Toderici, Lucas Theis, Nick Johnston, Eirikur Agustsson, Fabian Mentzer, Johannes Ballé, Wenzhe Shi, and Radu Timofte. *CLIC 2021: Challenge on Learned Image Compression*. 2021. [4](#)
- [28] Róbert Torfason, Fabian Mentzer, Eirikur Ágústsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Towards image understanding from deep compression without decoding. In *International Conference on Learning Representations*, 2018. [2](#)
- [29] G.K. Wallace. The jpeg still picture compression standard. *IEEE Transactions on Consumer Electronics*, 38(1):xviii–xxxiv, 1992. [2](#)
- [30] H. Wang, Yitong Wang, Zheng Zhou, Xing Ji, Zhifeng Li, Dihong Gong, Jingchao Zhou, and Wenyu Liu. Cosface: Large margin cosine loss for deep face recognition. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5265–5274, 2018. [5](#), [7](#)
- [31] Shurun Wang, Shiqi Wang, Xinfeng Zhang, Shanshe Wang, Siwei Ma, and Wen Gao. Scalable facial image compression with deep feature reconstruction. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 2691–2695. IEEE, 2019. [2](#)
- [32] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10039–10049, 2021. [2](#)
- [33] Z. Wang, E.P. Simoncelli, and A.C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thirtieth Seventh Asilomar Conference on Signals, Systems Computers, 2003*, volume 2, pages 1398–1402 Vol.2, 2003. [3](#), [5](#)