Identity preserving loss for learned image compression

Jiuhong Xiao, Lavisha Aggarwal, Prithviraj Banerjee, Manoj Aggarwal and Gerard Medioni

Introduction

- Face recognition system uses compression to reduce latency for the transfer between local devices and cloud server
- Existing standard compression methods (JPEG/HEVC) may deteriorate the recognition accuracy for compression artifacts on the face region

Contribution

- Our method achieves ~38% file size of HEVC results with similar image quality and recognition accuracy
- Plug-in compression model without retraining/finetuning downstream model
- Compression model is robust to downstream model change



 Models trained with IP loss show 0.15% lower FRR@FAR=1% and reduce 62% BPP than HEVC. • Our results show tradeoff between image quality, recognition performance and file size. We balance the tradeoff by hyperparameter setting (%) FRR@FAR=1% 1.6 1.4 1.2

[Legend explanation] JPEG/HEVC: Standard compression baselines, L2/MS-SSIM: Reconstruction loss function, IPR: IP loss + reconstruction(REC) loss



Experiments

- (Left) Domain-specific training reduces 29% BPP than training in general domain with same image quality (experiments on CelebA-HQ dataset)
 - (Right) After changing downstream model, model trained with IP loss decreases 0.76% FRR@FAR=1 and 62% BPP than HEVC



Conclusion





Compression model trained with identity preserving loss achieves

higher compression ratio than standard compression baselines

with undegraded recognition accuracy.

Contact

Author Contact xjiuhong@amazon.com Amazon Science www.amazon.science

This is for research purposes only.



