Does lossy image compression affect racial bias within face recognition?

Seyma Yucer, Matthew Poyser, Noura Al Moubayed and Toby P. Breckon

Durham University

Motivation

Over the decades companies, nonprofits, and governments have deployed **an increasing number of face recognition systems** across various areas to make autonomous decisions **for millions of users.** However, such wide-scale adoption within real-world scenarios heightens public concern about the **potential for abuse and the harmful effects of face recognition due to the presence of bias.**

This study examines whether lossy image compression adversely impacts phenotype-based racial bias within face recognition during training and testing.

Phenotype-based Racial Bias Analysis Methodology

We adopt a racial bias analysis methodology that uses facial phenotype attributes for face verification task [1].

Figure 1: Adapted Facial phenotype attributes and their categorisation from [1].

Lossy Compression & Subsampling

We investigate the effect of sampling ratio on racial phenotype grouping performances and **compare the default 4:2:0 subsampling with the 4:4:4 no chroma-subsampling factor,** which keeps luminance and colour information unchanged.

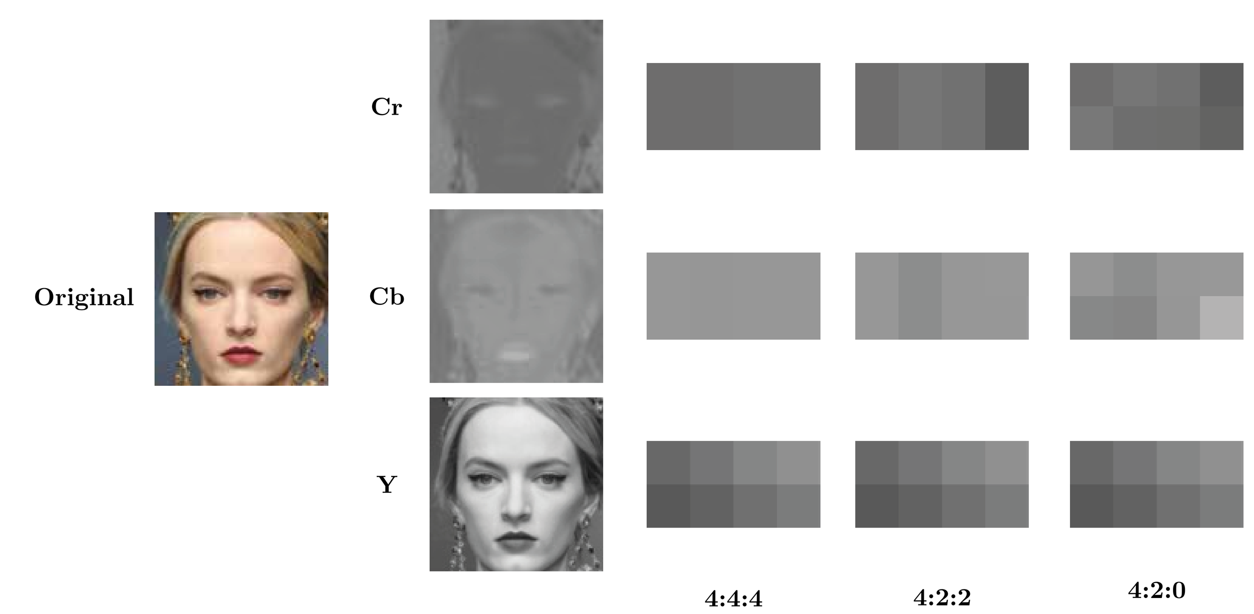

Figure 2: Chroma subsampling operation on different rates (4:2:0, 4:2:2, 4:4:4). Each rate differs according to how many pixels will be the same in the block.

Compression Level Selection

We analyse uniformly distributed compression levels on the RFW benchmark face recognition dataset using PSNR and down-selected **5, 10, 15, and 95 compression levels** in which quality decrease is most apparent.

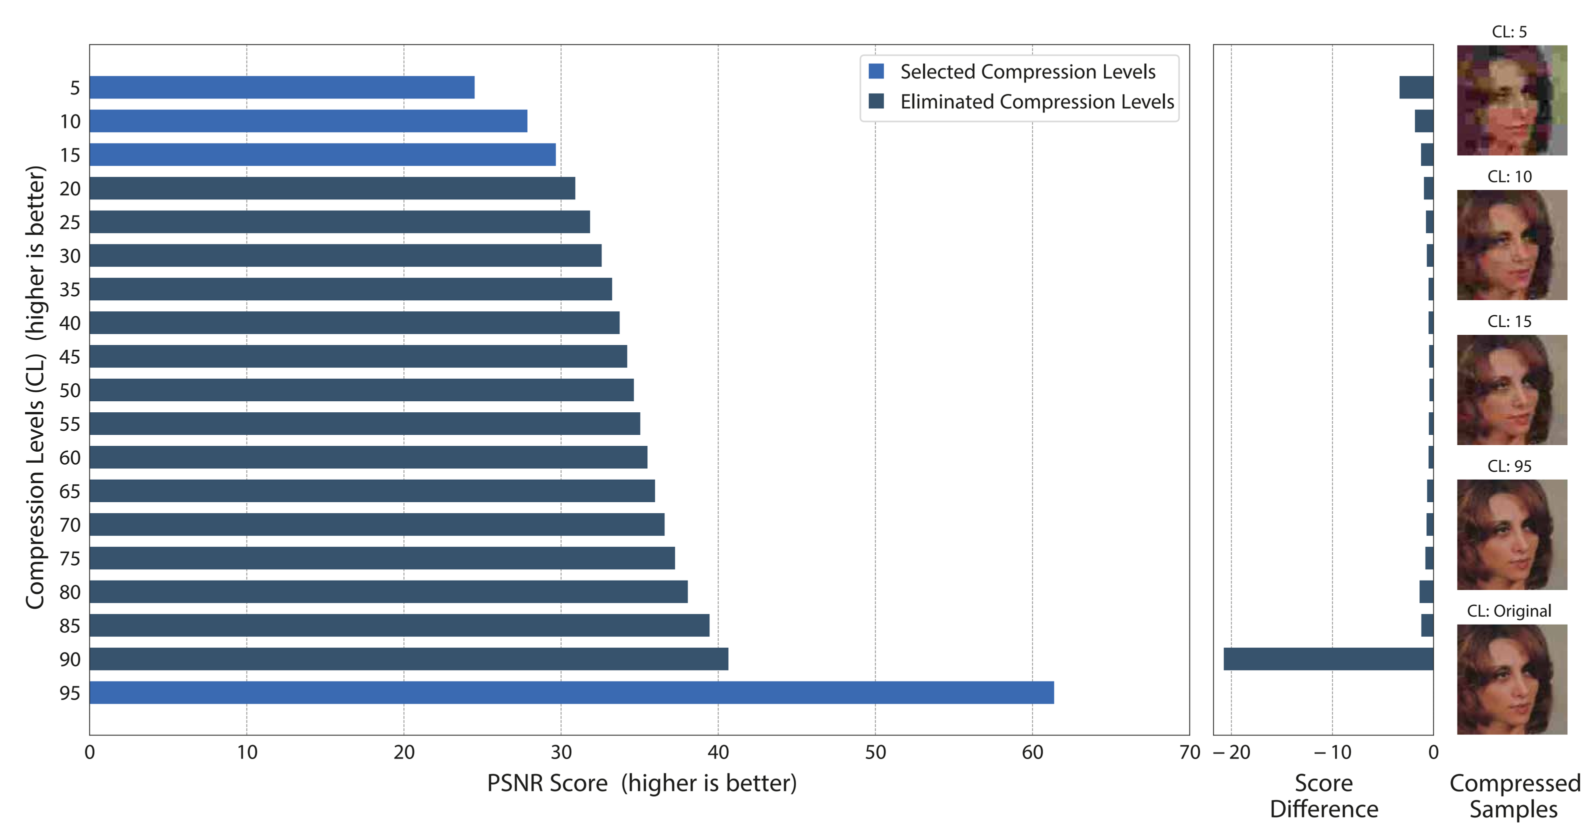

Figure 3: PSNR Scores of RFW dataset on different compression levels (CL). Relative score difference shows how much the quality changes at each level.

References

[1] S. Yucer, F. Tektas, N. Al Moubayed, and T. P. Breckon, Measuring hidden bias within face recognition via racial phenotypes in Proc. WACV, 2022

[2] Deng, Jiankang, et al. "Arcface: Additive angular margin loss for deep face recognition." Proc. CVPR, 2019

Results

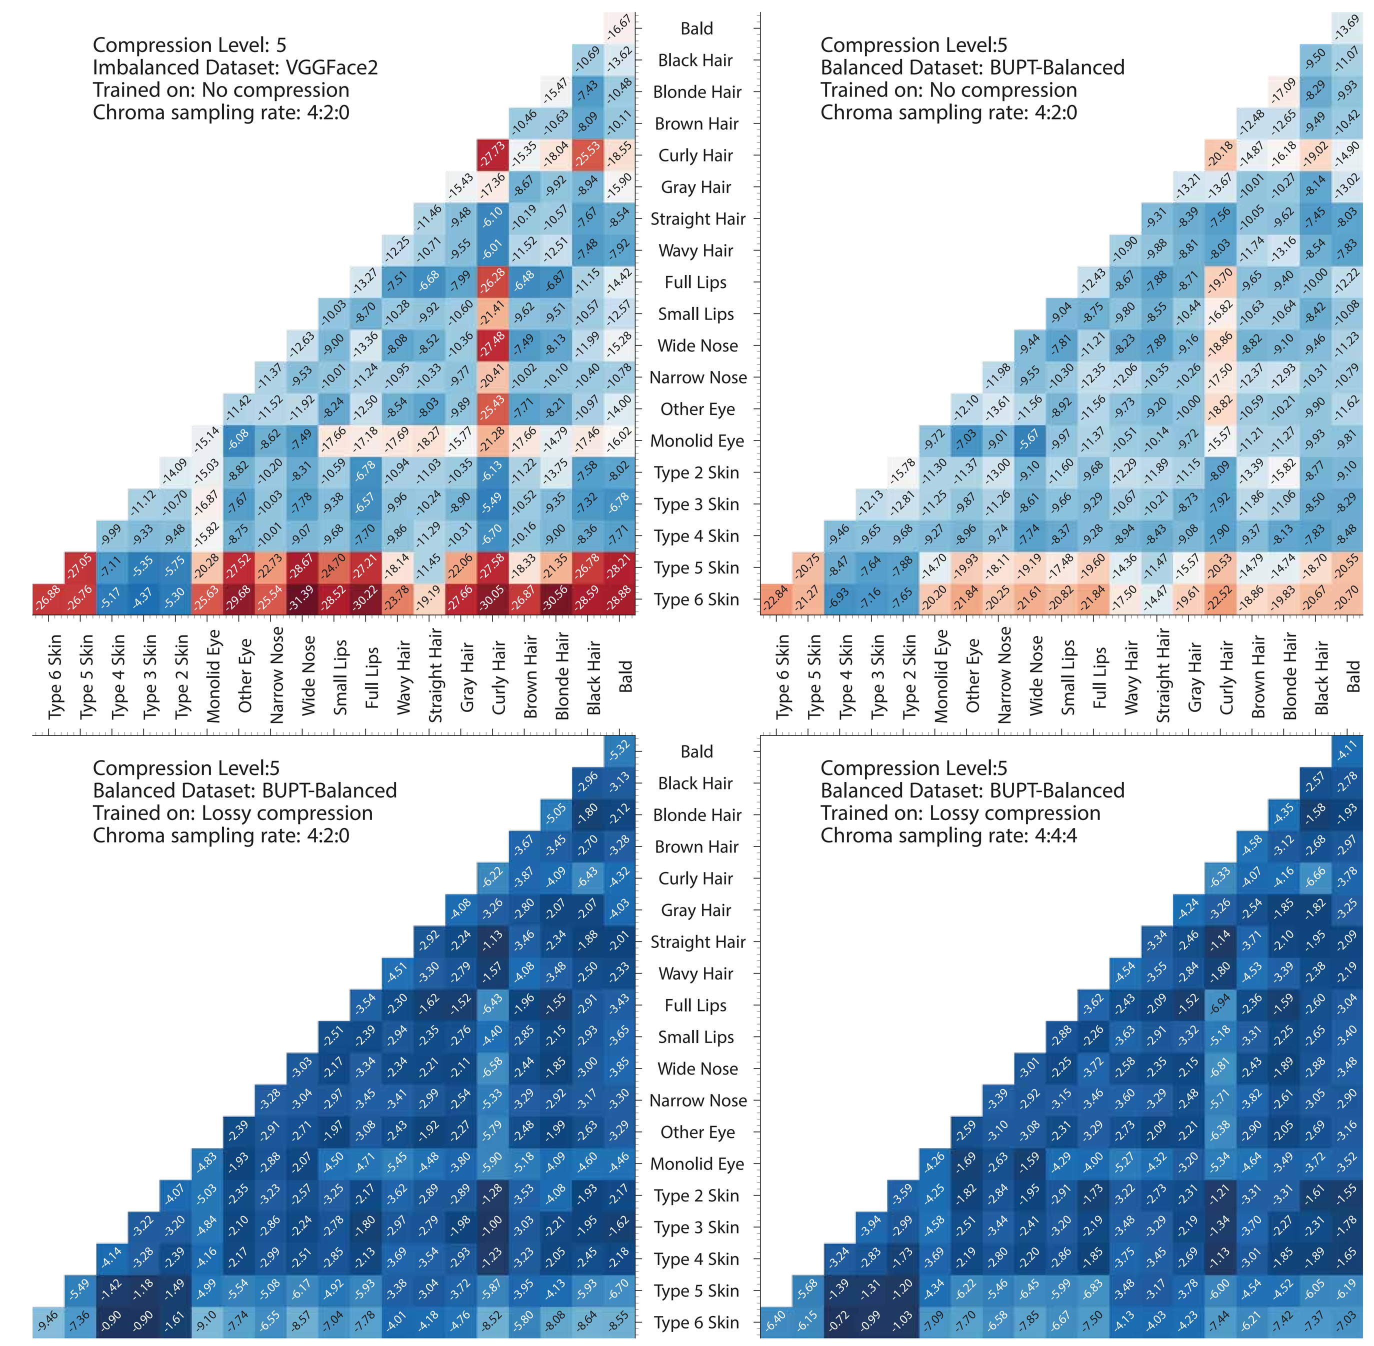

Figure 4: VGGFace2 and BUPT-Balanced training sets, compressed RFW test imagery (q=5); FMR performance differences of cross-attribute based pairings. Each cell depicts $FMR_{original} - FMR_{5}$.

Using ArcFace [2] as a baseline, we find:

- Compressed test imagery **increases FMR and decrease accuracy significantly on specific phenotypes, including dark skin tone, wide nose, curly hair, and monolid eye** across other phenotypes.

- The use of compressed imagery during training does make the models **more resilient and limits the performance degradation.**

- Removing chroma subsampling **improves FMR for specific phenotype categories more affected by lossy compression**.

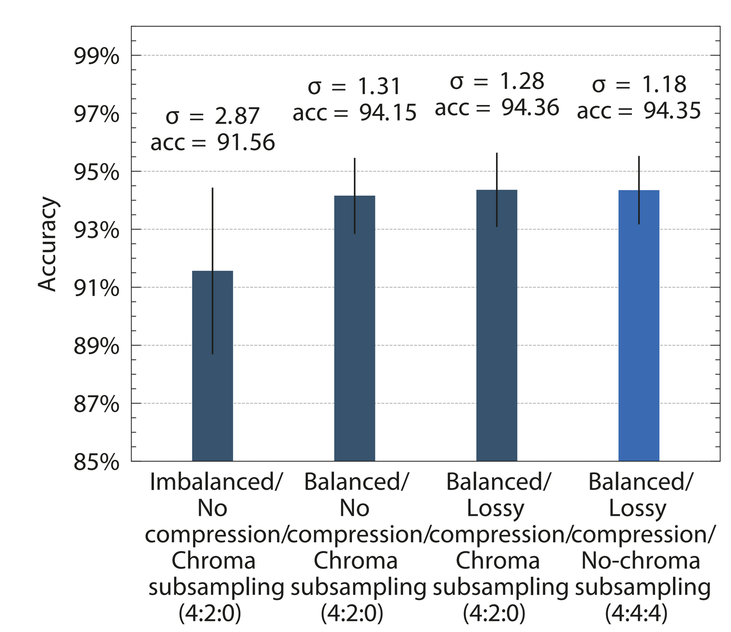

Figure 5: Accuracy and standard deviation of all attribute categories and their comparison on different training strategies using compressed (q = 75) of RFW.