


JPEG Compression Increases Facial Recognition Error Rate for Non-Caucasian Faces, Study Finds

Originally published at <https://www.unite.ai/jpeg-compression-increases-facial-recognition-error-rate-for-non-caucasian-faces-study-finds/>




Main image: DALL-E 2.

A new study from the UK has concluded that lossy compression techniques in JPEG images can have an adverse influence on the effectiveness of facial recognition systems, making such systems more likely to incorrectly identify a non-Caucasian person.

The paper states:

'Via an extensive experimental setup, we demonstrate that common lossy image compression approaches have a more pronounced negative impact on facial recognition performance for specific racial phenotype categories such as darker skin tones (by up to 34.55%).'

The results also indicate that [chroma subsampling](#), which reduces the color information (rather than the brightness information) across sections of a face image increases the False Matching Rate (FMR) across a range of tested datasets, many of which are standard repositories for computer vision.



Chroma subsampling operations on a source image, at varying rates, have a clear effect on the extent to which detail is preserved, and the extent to which sub-tones simply 'blend' into each other, sacrificing detail and determining features. Please note that this image in itself may be subject to compression, and refer to the source paper for accurate resolution. Source: <https://arxiv.org/pdf/2208.07613.pdf>

Chroma subsampling is applied as an additional economic measure in JPEG compression because people are less able to perceive reductions in the complexity and range of a color-bands than computer vision systems, which take these 'aggregations' far more literally than we do.

The researchers for the new study have found that removing chroma subsampling from the compression process lessens this negative effect by up to 15.95%, though it does not completely remove the problem.

The study also asserts that training on uncompressed (or less compressed) data will *not* resolve the problem if the inference-time images are compressed. Effectively, this means that training a facial recognition model on less-compressed imagery will not resolve the bias if the final production model is fed images that have the stated compression issues.

The authors report*:

'[The] use of lossy image compression during inference adversely affects the performance of contemporary [face recognition approaches](#) on a subset of race-related facial phenotype grouping (i.e. darker skin tones, monolid eye shape) and that its effect is present regardless of whether compressed imagery is used for model training.'

The paper underlines the consequences of image compression on the computer vision research sector, which were spelled out in some detail in a [2021 study](#) from the University of Maryland and Facebook AI.

It's a [difficult issue to remediate](#); even if the storage and bandwidth issues that make compression necessary were eliminated overnight, and even if all the low-quality images that populate twenty or more years of datasets in the sector were suddenly recompressed at a better rate from high-quality sources, it would represent a 'reset' of the continuity of academic benchmarking tools over the past few decades. The CV community has, in effect, [become accustomed](#) to the problem, to the point where it represents a notable technical debt.

Racial [bias](#) in facial recognition (FR) has [become](#) a [hot media topic](#) in recent years, prompting a concerted effort in the research community to eliminate it from affected systems. However, the dependence on the global research body on an [excessively limited](#) number of 'gold standard' datasets, many of which are either [not racially balanced](#) or [poorly labeled](#) in this respect, exacerbates the challenge.

The researchers of the new paper additionally note a dissonance between image acquisition standards and the standards set by the general run of facial recognition benchmarks, stating*:

'[Existing] image acquisition standards for face recognition systems such as [ISO/IEC 19794-5](#) and [ICAO 9303](#) propose both image-based (i.e. illumination, occlusion) and subject-based (i.e. pose, expression, accessories) quality standards to ensure facial image quality.

'Accordingly, facial images should also be stored using lossy image compression standards such as [JPEG](#) or JPEG2000; and identifiable for gender, eye colour, hair colour, expression, properties (i.e. glasses), pose angles (yaw, pitch, and roll), and landmark positions.

'However, common face recognition benchmarks do not conform to the ISO/IEC 19794-5 and ICAO 9303 standards. Moreover, in-the-wild samples are often obtained under the varying camera and environmental conditions to challenge the proposed solutions.

'Nevertheless, most facial image samples within such datasets are compressed via lossy JPEG compression.'


The authors of the new work state that their future efforts will examine the impact of lossy image quantization on diverse face recognition frameworks, and offer possible methods to improve the fairness of these systems.

The [new paper](#) is titled *Does lossy image compression affect racial bias within face recognition?*, and comes from three researchers at Imperial College London, together with one from the InsightFace deep face analysis [library](#).

Data and Method

For their experiments, the researchers used the [ImageMagick](#) and [libjpeg](#) open source libraries to create versions of the source data images at various increments of compression.

For an initial overview of the effects of compression, the authors studied the effects of Peak signal-to-noise ratio ([PSNR](#)) on four different levels of JPEG compression on the Racial Faces in-the-Wild ([RFW](#)) dataset.




Additionally, the researchers reproduced the effects of compressed and non-compressed training by removing chroma subsampling, in order to measure its effect on performance.

Results

The False Matching Rate (FMR) across these generated datasets were then studied. The criteria the researchers were looking for were predefined [phenotypes](#) relating to racial characteristics *Skin Type* (1, 2, 3, 4, 5 or 6), *Eyelid Type* (Monolid/Other), *Nose Shape* (Wide/Narrow), *Lip Shape* (Full/Small), *Hair Type* (Straight/Wavy/Curly/Bald), and *Hair Colour* – metrics drawn from the 2019 [paper](#) *Measuring Hidden Bias within Face Recognition via Racial Phenotypes*.

The paper states:

'We observe that for all down-selected compression levels $q = \{5, 10, 15, 95\}$, the FMR increases when additional lossy compression is applied, demonstrating that compression level 5 (the highest compression rate) results in the most significant decrease in FMR performance, whilst compression level 95 (the lowest compression rate) does not result in any noticeable FMR performance differences.'



A sample from the paper's extensive results charts, which are too large and numerous to reproduce here – please see the source paper for better resolution and find we see the gamut of FMR performance across increasingly degraded/compressed face images for VGGFace2, in a range that includes uncompressed or little-con

The paper concludes:

'Overall, our evaluation finds that using lossy compressed facial image samples at inference time decreases performance more significantly on specific phenotypes, including dark skin tone, wide nose, curly hair, and monolid eye across all other phenotypic features.'

'However, the use of compressed imagery during training does make the resulting models more resilient and limits the performance degradation encountered: lower performance amongst specific racially-aligned sub-groups remains. Additionally, removing chroma subsampling improves FMR for specific phenotype categories more affected by lossy compression.'

* My conversion of the authors' inline citations to hyperlinks.

First published 22nd August 2022.

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



Discover More