CLEANOR LABS · TECHNICAL REPORT · JUNE 2026

Iso-Quality Storage Efficiency of JPEG, WebP, and AVIF: A Reproducible Benchmark on Lossless and Real-World 12-Megapixel Photographs

Cleanor Labs Research Team

Cleanor Research Labs · Correspondence: hello@cleanor.app

Preprint · https://cleanor.app/blog/jpeg-webp-avif-photo-compression-savings-benchmark

ABSTRACT

Photographs are the largest single consumer of storage on most smartphones, so the image format and quality at which they are stored is the dominant lever for reclaiming space. We benchmark three formats — JPEG, WebP, and AVIF — under an *iso-quality* protocol that holds every codec to the same measured visual quality (SSIM) and compares the bytes required. Across two corpora — the 24-image lossless Kodak set and eight real 12-megapixel photographs — and an eight-step quality ladder (768 total encodes via `sharp`/libvips, scored with FFmpeg SSIM and PSNR), AVIF is the most efficient format, requiring 36–52% fewer bytes than JPEG at matched quality from a lossless source and ~29% fewer on real photographs. WebP is a consistent ~34% smaller than JPEG from a pristine original, but *inflates* by 25–57% when re-compressing the already-JPEG photographs typical of a phone. We also show that the nominal "quality" setting is not comparable across formats, and that AVIF's advantage narrows toward the visually-lossless limit. All code and the 768-encode dataset are released for reproduction.

Keywords: image compression, AVIF, WebP, JPEG, SSIM, rate–distortion, mobile storage, HEIC

1. Introduction

On almost every full phone, photos and videos are the single biggest storage consumer — typically a larger slice than apps, messages, and downloads combined. The practical consequence is that the *format and quality* in which pictures are stored is the real lever for reclaiming space, far more than uninstalling a handful of applications.

Yet the everyday advice to "just compress your photos" hides considerable nuance. The same image can be stored as JPEG, WebP, or AVIF, at any nominal quality setting, producing widely different file sizes — and different *actual* visual quality even when the slider reads the same number. A poor choice can shrink a photo with no perceptual benefit, retain quality the viewer never needed, or — as we show for WebP — make a photograph *larger* than it began. We therefore ask a deliberately concrete question: *for the same picture, viewed at the same quality, how many bytes does each format actually need?* That number determines whether re-compressing a library is worthwhile, and which format to use.

2. Method

2.1 Corpora

A single test set cannot answer both "which format is fundamentally more efficient?" and "what happens when I re-compress the photos already on my phone?" We therefore use two corpora. **Corpus 1** is the *Kodak True Color* set [2]: 24 lossless PNG images at 768×512 (or 512×768), the standard reference corpus in image-compression research, spanning portraits, landscapes, architecture, and fine textures. Because the source is lossless, it is a fair test of each format's intrinsic efficiency. **Corpus 2** is eight real photographs at 12.2 megapixels (4032×3024) — the native resolution of a modern phone camera — supplied as already-compressed JPEGs (mean source size 1.54 MiB). This corpus measures the real-world operation a cleanup tool performs: re-compressing an existing JPEG rather than encoding a pristine master.

2.2 Encoding

Every image was encoded to all three formats across an eight-step quality ladder (40, 50, 60, 70, 80, 90, 95, 98), for 768 encodes total (Kodak: 24×3×8 = 576; 12 MP: 8×3×8 = 192). All encoding used `sharp` 0.35.1 / libvips 8.18.3 [8] — the same engine behind production web tooling — via libjpeg-turbo (baseline JPEG), libwebp (WebP), and libaom (AVIF, AV1) at default effort 4.

2.3 Quality metrics

Each candidate was scored against its own original with FFmpeg [7], recording two metrics. **SSIM** (structural similarity [1]) compares image structure — edges, textures, and local patterns — the way a human observer does, on a 0–1 scale where 1.0 is identical; we use the combined "All" channel. As a rule of thumb, SSIM ≥ 0.98 is "visually lossless" for normal viewing. **PSNR** (peak signal-to-noise ratio, dB) is a traditional error measure recorded as a cross-check; it agreed with SSIM on the ranking of the formats at every quality level (e.g., at quality 80: AVIF 44.1 dB > WebP 40.8 dB > JPEG 38.3 dB), so SSIM-based figures are reported throughout.

2.4 Iso-quality comparison

Formats cannot be compared by giving each the same "quality 80," because that number means something different in every codec. We therefore use an *iso-quality* protocol: for each image and format we interpolate the file size needed to reach a fixed SSIM target, then average across the corpus. Every format is thus held to the same measured visual quality, and the only thing that varies is the bytes required — precisely the quantity relevant to storage.

2.5 Reproduction

The study re-runs from free, open tools. (i) Obtain a test set — the lossless Kodak images, or one's own photographs. (ii) Encode each to JPEG, WebP, and AVIF across quality 40–98 with any libvips/`sharp`, ImageMagick, or standalone `cjpeg`/`cwebp`/`avifenc` build, recording output sizes. (iii) Score each result against its original with `ffmpeg -i orig -i cand -lavfi ssim -f null -` (and `psnr`). (iv) For each format, find the size that reaches a fixed SSIM and compare. Our exact scripts and the 768-encode dataset accompany this report.

3. Results

3.1 Format efficiency on lossless sources

On the lossless Kodak set, AVIF is the most efficient format at the quality at which most photographs are stored, with WebP second and JPEG least efficient. At "good" quality (SSIM 0.95) one image requires 55.5 KB as JPEG, 36.5 KB as WebP (−34%), and 26.7 KB as AVIF (−52%). AVIF's lead is largest at everyday quality and narrows toward the visually-lossless limit; at SSIM 0.99 WebP's steady −34% actually surpasses AVIF (Table 1).

Table 1. Iso-quality byte cost relative to JPEG, lossless Kodak set (n = 24).

Visual quality (SSIM)	AVIF vs JPEG	WebP vs JPEG
0.95 (good)	−52%	−34%
0.98 (high)	−36%	−34%
0.99 (near-lossless)	−17%	−34%

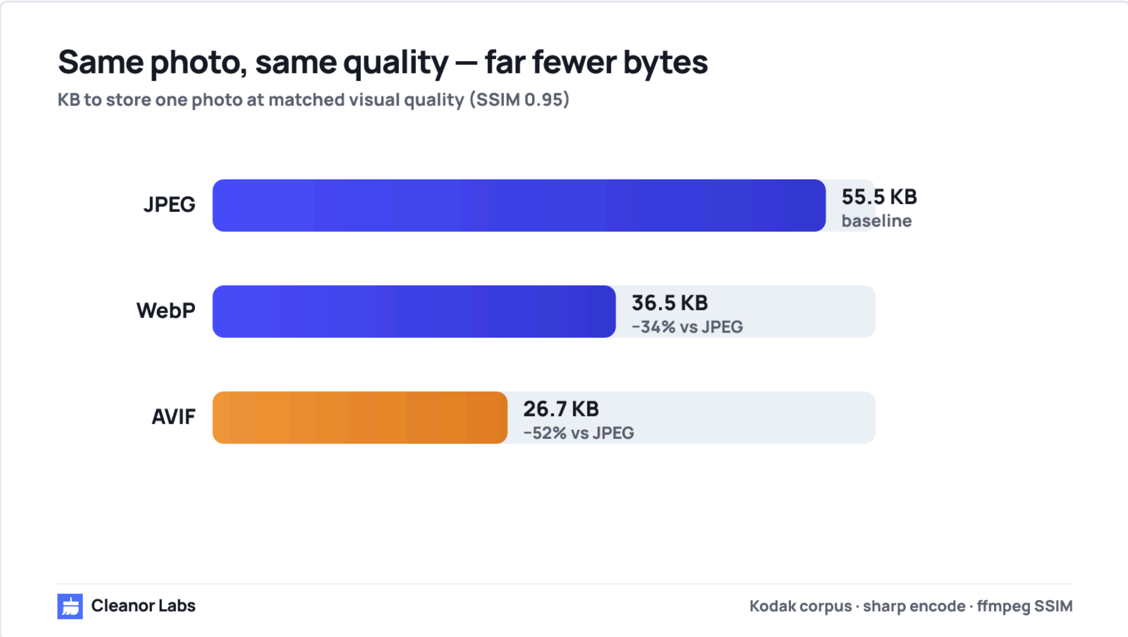

Figure 1. Bytes required to store one Kodak image at matched visual quality (SSIM 0.95). AVIF needs roughly half the bytes of JPEG; WebP about a third fewer.

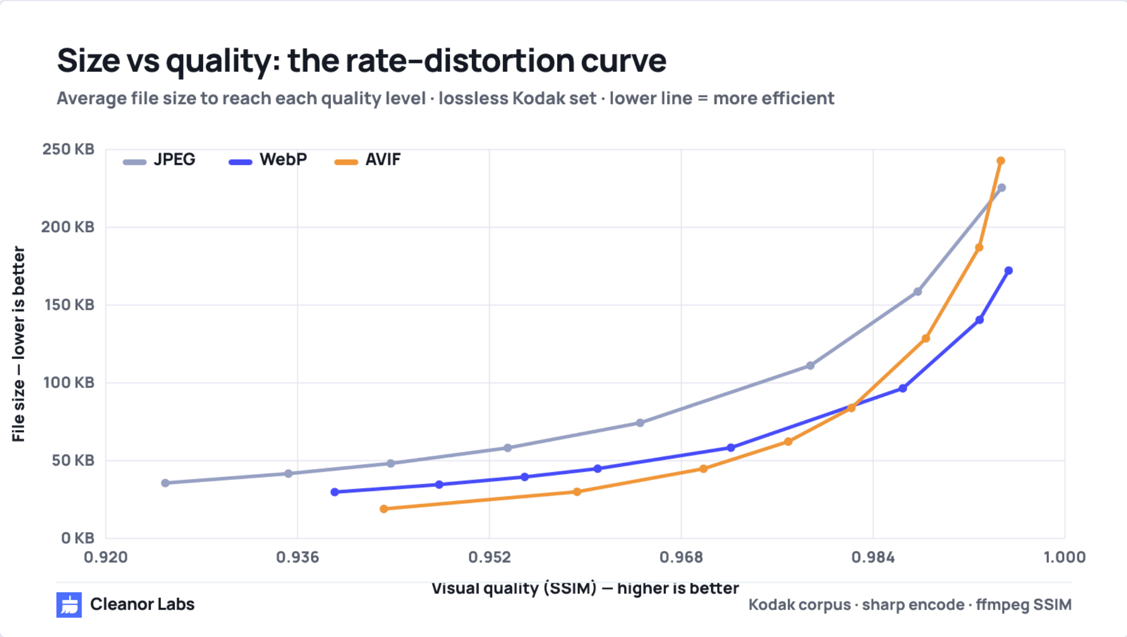

Figure 2. Rate–distortion curves (mean file size vs. achieved SSIM) on the lossless Kodak set. A lower line is more efficient; AVIF (orange) is lowest across most of the range, with WebP (blue) overtaking it only near the lossless limit.

3.2 Real-world re-compression at 12 megapixels

Lossless sources flatter every codec, so we repeated the protocol on eight real 12-megapixel JPEGs, re-compressing each to its original quality (SSIM 0.98). AVIF retained a ~29% advantage over an equal-quality JPEG, but WebP produced *larger* files than JPEG — by 25% at high quality and up to 57% near lossless (Table 2). At SSIM 0.98, the mean per-photo size was 1.35 MiB (JPEG), 1.69 MiB (WebP), and 0.96 MiB (AVIF).

Table 2. Iso-quality byte cost relative to JPEG, real 12-megapixel photographs (n = 8).

Visual quality (SSIM)	AVIF vs JPEG	WebP vs JPEG
0.95	−41%	−13%
0.98	−29%	+25%
0.99	−22%	+57%

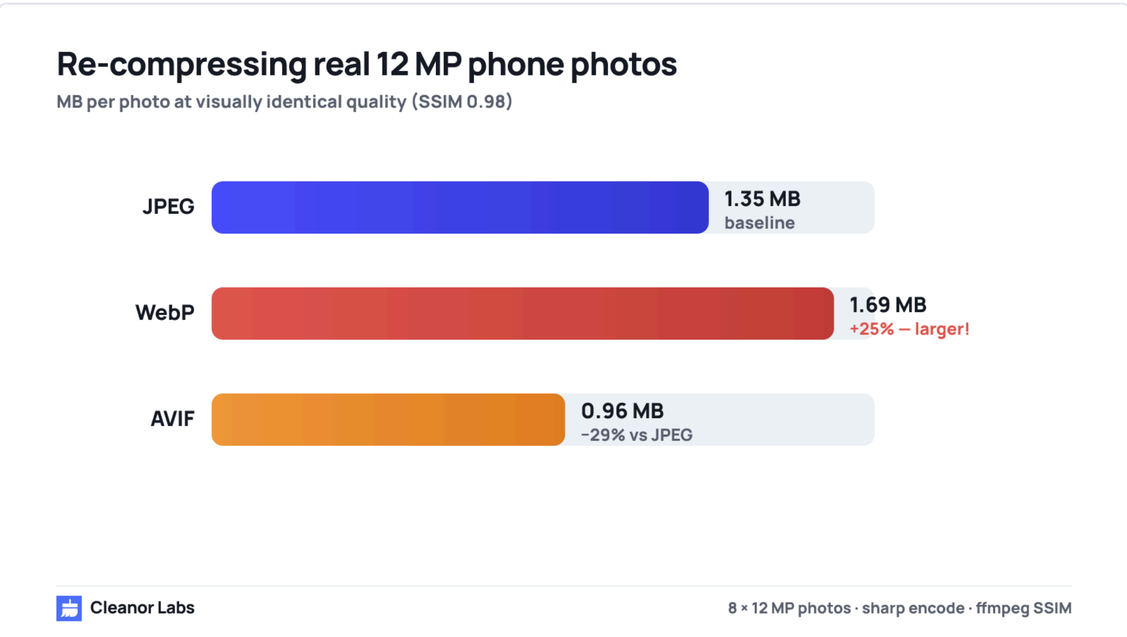

Figure 3. Mean size per 12-megapixel photograph at matched quality (SSIM 0.98). AVIF is smaller than JPEG; WebP is larger.

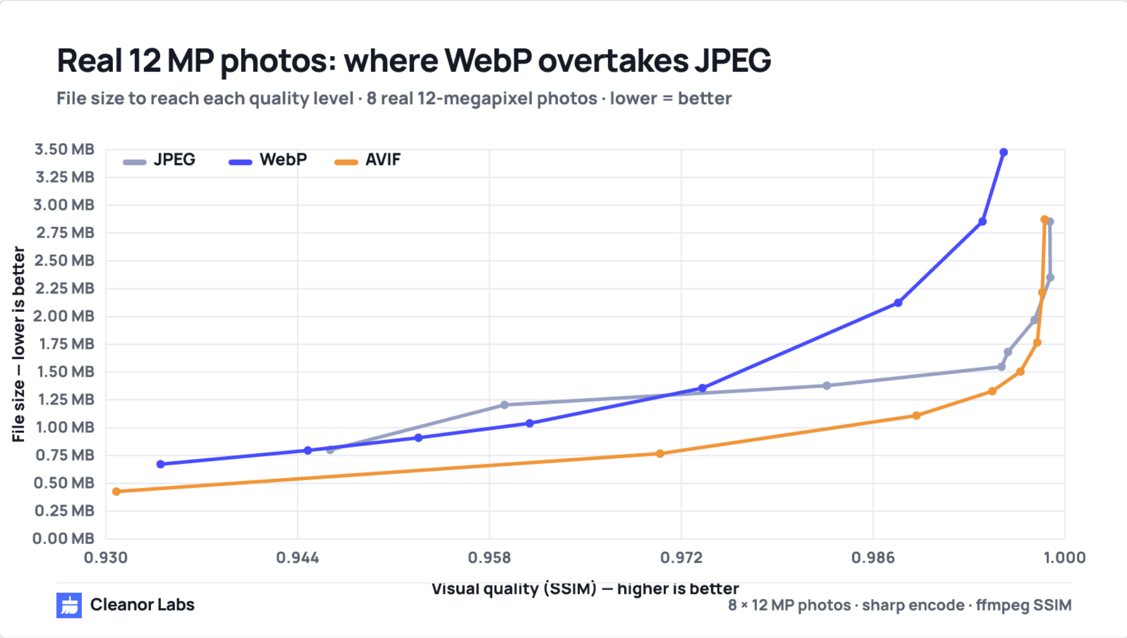

Figure 4. Rate–distortion curves on real 12-megapixel photographs. AVIF remains lowest, while the WebP line climbs above JPEG past the high-quality mark — re-compressing existing phone JPEGs to WebP enlarges them.

Two effects explain WebP's inflation. First, the source photographs are already JPEG, and re-encoding JPEG-to-JPEG is unusually cheap because the transform basis is reused. Second, libwebp is inefficient at high fidelity on large images. Together, the starting JPEG is a low bar that WebP cannot beat, whereas AVIF's encoder is not subject to either effect.

3.3 The nominal quality setting is not comparable

At quality = 80 on the Kodak set, the three formats produced different sizes at different actual quality: JPEG 74.3 KB at SSIM 0.965, WebP 58.4 KB at 0.972, and AVIF 83.8 KB at 0.982. AVIF's file is the largest here only because it delivered the highest quality. Compared on the slider the result is meaningless; compared at matched measured quality, AVIF is smallest. Table 3 gives the full curve.

Table 3. Mean Kodak file size (KB) and achieved SSIM at each nominal quality setting.

Quality	JPEG (KB / SSIM)	WebP (KB / SSIM)	AVIF (KB / SSIM)
40	35.7 / 0.925	29.8 / 0.939	19.0 / 0.943
60	48.2 / 0.944	39.5 / 0.955	44.8 / 0.970
80	74.3 / 0.965	58.4 / 0.972	83.8 / 0.982
90	111.1 / 0.979	96.5 / 0.986	128.7 / 0.988
98	225.5 / 0.995	172.2 / 0.995	242.8 / 0.995

4. Discussion

AVIF saves space in every scenario we tested — pristine exports and re-compressed phone photographs alike — making it the safe default for shrinking a photo library: format-for-format at matched quality, a 12-megapixel image needs about 1.0 MiB as AVIF versus ~1.4 MiB as JPEG, and is smaller still relative to the ~1.6 MiB original on the device. WebP is an excellent web format *from clean originals* but a poor choice for re-compressing existing high-quality photographs. Practically, users should target a visual result (SSIM ≈ 0.98) rather than a slider value, since the same number yields different sizes and quality across formats. One caveat for iPhone owners: HEIC, the HEVC-based format Apple captures by default, is already of broadly AVIF-class efficiency [6], so the large gains from re-compression come from older or shared JPEGs, not from the newest shots.

5. Limitations

SSIM is a single perceptual metric; PSNR was recorded as a cross-check and agreed, but all headline figures are SSIM-based. The Kodak set is lower-resolution — its efficiency *percentages* are resolution-independent, but absolute sizes scale with megapixels, which motivated the separate 12-megapixel corpus. On those real JPEGs the highest settings sit near the top of the measurable SSIM range, so the near-lossless (0.99) figures are the least precise, while the 0.95 and 0.98 results are firm. AVIF here is libaom via libvips at default effort 4; a slower effort could reduce AVIF sizes further. All 12-megapixel sizes are reported in MiB (1024^2 bytes).

6. Conclusion

Under a fair, matched-quality comparison, AVIF is the most storage-efficient of the three formats for photographs, saving roughly a third to a half over JPEG depending on the target quality, and remaining the only format that reliably shrinks the JPEGs already on a phone. WebP's gains are real only from pristine originals and can reverse on real-world re-compression. The full method and dataset are released so that any reader can reproduce these numbers.

References

1. Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, 2004.
2. Eastman Kodak Company, "Kodak Lossless True Color Image Suite (PhotoCD PCD0992)," 24-image reference set.
3. Alliance for Open Media, "AV1 Bitstream & Decoding Process Specification," 2019; AVIF Image File Format.
4. Google, "WebP: A new image format for the Web," developer documentation.
5. ITU-T Rec. T.81 | ISO/IEC 10918-1, "Digital compression and coding of continuous-tone still images (JPEG)."
6. ISO/IEC 23008-12, "High Efficiency Image File Format (HEIF)"; ISO/IEC 23008-2 (HEVC).
7. FFmpeg project, "FFmpeg multimedia framework," ssim and psnr filters. https://ffmpeg.org
8. J. Cupitt and K. Martinez, "libvips: a fast image processing library"; `sharp` Node.js binding. https://sharp.pixelplumbing.com
9. Cleanor Labs, "The Geography of Storage Anxiety: 117,000 Searches Mapped," 2026. cleanor.app/blog/storage-anxiety-by-country-search-data-2026
10. Cleanor Labs, "What Fills Up Phone Storage in 2026," 2026. cleanor.app/blog/phone-storage-problems-2026-search-demand-study

How to cite: Cleanor Labs Research Team. "Iso-Quality Storage Efficiency of JPEG, WebP, and AVIF: A Reproducible Benchmark on Lossless and Real-World 12-Megapixel Photographs." Cleanor Labs Technical Report, June 2026. https://cleanor.app/blog/jpeg-webp-avif-photo-compression-savings-benchmark