

# Video Quality Metrics Cheat Sheet

PSNR, SSIM, MS-SSIM, VMAF and the 6-point JND rule. Updated 2026-05-16.


## Metric overview

Metric	Scale	What it measures	When to use
<b>PSNR</b>	30-50 dB	Average pixel-by-pixel error in decibels. Sanity check, historical comparison.	
<b>SSIM</b>	0 - 1	Local luminance + contrast + structure. Image work, legacy codec papers.	
<b>MS-SSIM</b>	0 - 1	SSIM at 5 scales, weighted sum.	Codec research, offline UHD eval.
<b>VMAF default</b>	0 - 100	ML fusion of VIF, ADM, motion (1080p SDH).	Streaming ladder tuning.
<b>VMAF NEG</b>	0 - 100	VMAF without enhancement-gain inflation.	Fair codec A vs B comparisons.
<b>VMAF phone</b>	0 - 100	VMAF recalibrated for phone viewing distance.	Mobile-only encoding ladders.
<b>VMAF 4K</b>	0 - 100	VMAF for 4K TV at 1.5 picture heights.	Premium UHD content.
<b>XPSNR</b>	30-50 dB	Perceptually weighted PSNR (Fraunhofer).	V/C encoding, HDR, low-cost monitor.
<b>CAMBI</b>	0 - 24	Banding-only score (higher = more banding).	HDR pipeline, gradient-heavy content.

## VMAF score scale and the six-point rule

0 - 100 perceptual scale. 6 VMAF points = 1 Just Noticeable Difference (75% of viewers notice).

Keep adjacent ABR rungs within 2 VMAF points so quality switches are invisible.



## FFmpeg one-liners

Compute VMAF + PSNR + SSIM + MS-SSIM at once (default model):

```
ffmpeg -i enc.mp4 -i ref.mp4 -lavfi \
  "libvmaf=feature='name=psnr|name=float_ssim|name=float_ms_ssim':log_fmt=json:log_path=m.json" \
  -f null -
```

Conform a 720p encode against a 1080p reference before measuring:

```
ffmpeg -i enc_720p.mp4 -i ref_1080p.mp4 -lavfi \
  "[0:v]scale=1920:1080:flags=bicubic[d];[d][1:v]libvmaf" -f null -
```

GPU-accelerated on NVIDIA (FFmpeg 6.1+, ~6x faster at 1080p/4K):

```
ffmpeg -hwaccel cuda -i enc.mp4 -i ref.mp4 -lavfi libvmaf_cuda -f null -
```

## Production rules of thumb

- Track the 5th-percentile, not just the mean - one bad scene wrecks the audience experience.
- State the VMAF model in every report - default, NEG, phone, and 4K can differ by 10 points.
- Use NEG for codec comparisons, default for pipeline evaluation.
- Never make a codec choice on PSNR alone when bitrates differ by more than 10%.
- Conform resolution and frame rate before measuring - scaling itself affects PSNR/SSIM.
- Match the metric model to the viewing device: phone for phones, 4K for UHD.
- Two adjacent ABR rungs should never differ by more than 2 VMAF points.
- When metrics disagree, that disagreement is a signal: PSNR high + VMAF low = blurring.