Entropy coder selection — cheat sheet

CAVLC, CABAC, AV1 range coder, VVC multi-rate estimator — one page

Companion to article 4.10 of Block 4. Print on A4. (c) 2026 Fora Soft. v1.0

1. Entropy coder per codec

Codec	Coder	Alphabet	Contexts	Adaptation	Saving vs prior
MPEG-2 (1995)	Static Huffman	per-symbol VLC	none	none	baseline
H.264 Baseline	CAVLC	per-symbol VLC	5 tables	table choice	~5% over MPEG-2
H.264 High	CABAC	binary bins	~460	64-state per-slice	~10-15% over CAVLC
HEVC (2013)	CABAC only	binary bins	~500	64-state per-slice	~6-8% over H.264 CABAC
AV1 (2018)	Range coder	up to 16/symbol	~1000	per-symbol	~3-5% over HEVC
VVC (2020)	CABAC	binary bins	~700	multi-rate	~2% over HEVC

2. CAVLC vs CABAC at a glance

CAVLC - five short steps: coeff_token, trailing-ones signs, levels, total_zeros, run_before.

- Whole-bit codewords, table choice depends on what was just coded.
- Cheap to parse: a few cycles per coefficient block. Parallelises well.

CABAC - three stages: binarisation, context modelling, binary arithmetic coding.

- Fractional bits per symbol; within 1% of the entropy limit.
- Serial throughput: one bin per 10-20 cycles; main bottleneck in 4K hardware decoders.

3. AV1 range coder — what changed after CABAC

- Coder borrowed from Daala (Mozilla/xiph), arxiv.org/abs/1608.01947.
- Multi-symbol: up to 16 alphabet values coded in one engine call (no binarisation).
- Per-symbol probability adaptation - the table updates after every coded symbol.
- Multiple frame_context_idx tables in memory; pick one per frame to track scene shifts.
- Equivalent to ~4 binary bins per call. Hardware: fewer cycles per pixel than CABAC.
- Trade-off: more on-chip memory for context tables. Slower hardware AV1 rollout.

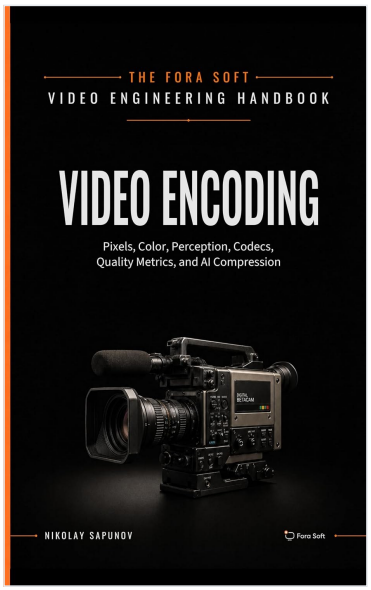
4. VVC — the small but real CABAC tweaks

- Multi-rate probability estimator - two parallel rate tracks per context.
- Estimator weights adapted to the current QP; tighter probability tracking at high bitrate.
- Context modelling for transform coefficients coupled to dependent quantization state.
- Net entropy-stage saving over HEVC: roughly 2% bitrate at the same quality.

5. Live-stream debugging quick reference

- Blocky garbage from random tune-in point to slice end (HEVC): CABAC re-init missing at slice.
- First few frames after seek decode wrong (H.264 High): cabac_init_idc not honoured.
- AV1 stream OK in dav1d, broken in software X: frame_context_idx tables not all implemented.
- Slow decode under VVC at low QP: dependent-Q forces sequential context updates.
- Encoder reports 5% larger files than spec budget: check that High profile (CABAC) is on, not Baseline.

Companion to: 'Entropy Coding in Detail: CAVLC, CABAC, and Arithmetic Coding'
www.forasoft.com/learn/video-encoding/articles/entropy-coding-cavlc-cabac



GO DEEPER · VOLUME 1 OF 7

Video Encoding: A Complete Guide to Video Compression

From Pixels and Color to H.264, HEVC, AV1, Quality Metrics, and AI

By Nikolay Sapunov, CEO at Fora Soft · 1,522 pages · Kindle & paperback on Amazon

Read it on Amazon → amazon.com/dp/B0HF52B65R

This sheet compares CAVLC, CABAC, the AV1 range coder and VVC's estimator and gives you the debugging hints for live pipelines. Volume 1 covers entropy coding as part of the whole encoder: how transforms and quantization feed the coder, context modelling, the throughput cost in hardware decoders, and the profile choices that switch it on.

WHAT YOU WILL FIND INSIDE

- Transform, quantization and entropy coding as one chain
- AV1 multi-symbol coding and its hardware trade-offs
- Context modelling and arithmetic coding in H.264 and HEVC
- Profiles that switch CAVLC or CABAC on and off

THE FULL HANDBOOK · 7 VOLUMES · 9,000+ PAGES · ALL ON AMAZON

Vol. 1 · Video Encoding — pixels and color to H.264, HEVC, AV1, quality metrics and AI
Vol. 2 · Video Streaming I — network transport, contribution and ingest, HLS, MPEG-DASH, CMAF
Vol. 3 · Video Streaming II — ABR, CDNs, players, WebRTC, DRM and the economics of delivery
Vol. 4 · Audio for Video I — PCM, codecs, loudness, accessibility, real-time voice processing
Vol. 5 · Audio for Video II — jitter buffers, lip sync, Atmos, quality metrics, neural codecs
Vol. 6 · AI for Video I — models and inference, computer vision, speech, search, generative video
Vol. 7 · AI for Video II — real-time pipelines, agents, serving, cost, compliance, verticals

Seeing corrupt frames or slow decode in a live HEVC or AV1 pipeline?

We will trace the bitstream, find the entropy-coding or profile mismatch behind the artefacts and fix the encoder or player configuration. 250+ video products shipped since 2005.

Book a 30-min call → **WhatsApp →** **Email us →**