

Neural ABR Deployment Decision Sheet

Pensieve · Comyco · Kairos · ABRL · Fugu · SODA — when neural ABR ships and when it does not

Four deployment shapes where neural ABR wins

Shape	When it wins
Large trace corpus	Operator has hundreds of thousands of real-user session traces (Netflix, YouTube, Disney+, Meta, Amazon, large MVPDs). Smaller operators rarely meet the data bar without a vendor.
Characterisable network	Wired residential broadband, fixed wireless, characterised mobile networks. Heterogeneous mobility (satellite, 4G↔5G handoffs, public Wi-Fi at venues) generalises poorly out of the box.
Switch suppression is tracked	Apps where the product team measures switches per minute as a brand quality benefit disproportionately — Comyco, Kairos, SODA all suppress switches more aggressively than simple rules.
Team to maintain the model	Production deployments need a retraining pipeline, A/B infrastructure that detects single-digit-percent shifts, observability tooling, and a deterministic fallback. Without those, lifecycle cost dwarfs QoE gain.

Four production failure modes


Failure mode	Symptom & fix
Trace distribution mismatch	Policy trained on one network distribution running on another. Pensieve underperformed Fugu in the Puffer trial for exactly this reason. Fix: retrain in situ on the deployment's own traces.
Training cost & operational drag	RL training is brittle; reward shaping is an art; fresh corpora can require thousands of GPU-hours to retrain. Imitation learning (Comyco, Fugu) is dramatically cheaper but still needs an expert and a maintained pipeline.
Opacity in production debugging	A neural policy is a black box. Build observability tooling — saliency maps, decision logging, counterfactual replays — before deployment, not after the first customer escalation.
Reward-function misalignment	A policy maximises the reward you train it on. Engagement-shaped rewards (switches, TTFF, time-to-1080p) beat QoE-shaped rewards in production. SODA's case study makes this argument explicitly.

Five tuning levers that actually move QoE

Lever	How to set it
Training method	RL when no expert is available (Pensieve, ABRL). Imitation learning when an offline expert exists (Comyco, Fugu). Imitation is the 2026 default.
Where the model sits	Pure policy (Pensieve, ABRL, Comyco) when team can absorb opacity. Learned predictor inside an MPC controller (Kairos, Fugu) for production debugging. Hybrid framing is the 2026 default.
Reward shaping	Academic default: quality – λ s-switches – λ r-rebuffer. Production: add TTFF penalty, time-to-1080p penalty, viewing-duration term. Retrain when weights change.
Retraining cadence	Monthly for stable wired networks; weekly for fast-evolving mobile; continuous lifelong learning (L-Comyco) where the pipeline supports it.
Fallback strategy	Mandatory deterministic fallback to BOLA or MPC when the policy returns a pathological action. Log the incident, A/B detect the drift, retrain or roll back.

Three production deployments documented at scale

Deployment	Architecture	Production result
Facebook ABRL (2020)	RL policy. 30M+ sessions/week.	+1.6% avg bitrate · –0.4% stalls. arXiv 2008.12858.
Stanford Puffer / Fugu (NSDI 2020)	MPC + in-situ learned predictor. 8,131 hours · 3,719 users.	0.13% stall ratio · beats Pensieve in the wild. nsdi20-paper-yan.
Amazon Prime Video SODA (SIGCOMM 2024)	Smoothness-optimised DP controller with QoE guarantees.	–88.8% switches · +5.91% view duration on production fleet.



GO DEEPER · VOLUME 3 OF 7

Video Streaming: A Complete Guide to Video Delivery

From ABR and CDNs to Players, WebRTC, DRM, and the Economics of Streaming

By Nikolay Sapunov, CEO at Fora Soft · 1,182 pages · Kindle & paperback on Amazon

[Read it on Amazon →](#)
amazon.com/dp/B0HH76ZYWC

This sheet shows where neural ABR earns its cost, how it fails in production and which deployments proved it at scale. Volume 3 covers adaptive bitrate in depth: bitrate ladders, buffer-based, throughput-based and hybrid rung selection, player stall recovery, and QoE monitoring.

WHAT YOU WILL FIND INSIDE

- Buffer-based versus throughput-based ABR
- Hybrid rung-selection algorithms such as MPC
- Player stall recovery and switch suppression
- QoE metrics to judge an ABR change

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

Deciding whether neural ABR is worth it for your player?

We will benchmark your current ABR on real session traces and tell you whether tuning BOLA or MPC beats a learned policy for your audience. 250+ video products shipped since 2005.

[Book a 30-min call →](#)
[WhatsApp →](#)
[Email us →](#)