


The AI Hook Map for Video Products

Latency tightens toward the camera; cost-per-minute drops toward the archive.

1 · Capture	2 · Real-time transport	3 · Server-side	4 · Post-processing	5 · Long-form analytics
user device	SFU / live path	during the session	queue, just after	over the archive
~ 16 to 33 ms latency budget	~ 100 to 200 ms latency budget	~ 3 to 10 seconds latency budget	~ 1 to 10 minutes latency budget	hours to days latency budget
~ \$0 per minute cost shape	\$ per active stream cost shape	\$ per active session cost shape	cheap per minute cost shape	spiky, periodic cost shape
Typical capabilities Background blur Noise suppression Beauty filters On-device denoise	Typical capabilities Live captions Live translation Live moderation In-call AI agent	Typical capabilities Action items Live analytics In-call OCR Live sentiment	Typical capabilities Meeting summary Chapter markers Highlight reels Archive denoise	Typical capabilities Video search Corpus moderation Training-data prep Vector indexing

Decision rule

Result needed before next frame? Hook 1. Inside a live call? Hook 2 or 3. Within minutes of end? Hook 4. Over the archive? Hook 5.



GO DEEPER · VOLUME 6 OF 7

AI for Video: A Complete Guide to Intelligence in Video Products

From Models and Inference to Computer Vision, Speech, Search, and Generative Video

By Nikolay Sapunov, CEO at Fora Soft · 1,444 pages · Kindle & paperback on Amazon

[Read it on Amazon →](#)
amazon.com/dp/B0HH6TKPZ1

This map shows the five places AI hooks into a video product, from capture to the archive. Volume 6 explains each capability you would hang there: computer vision, speech recognition and synthesis, search and RAG over video, generative video, and the deployment topologies that decide latency and cost.

WHAT YOU WILL FIND INSIDE

- Detection, tracking, segmentation and OCR
- ASR, TTS and diarization
- Search and RAG over a video archive
- Deployment topologies and their latency

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

Looking for the right place to add AI to your video product?

We will walk your product through the five hooks and come back with a short list of features and their cost shape. 250+ video products shipped since 2005.

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