Learning-AI Build-vs-Buy Worksheet

Run this gate for each AI feature before you commit. The drivers, not the model leaderboard, pick the lane. Convert every option to cost per learner per month before deciding.

A. The feature and the option

- [ ] We named the AI feature and what learner data it touches
- [ ] We considered all three options: buy a product, call a managed API, self-host an open model
- [ ] For a new or unproven feature, we start on a managed API to validate before optimizing

B. The four drivers

- [ ] Usage volume: learners, minutes, or messages per month is estimated
- [ ] Accuracy bar: how wrong the feature is allowed to be is defined
- [ ] Data sensitivity: whether the learner data can legally leave our walls is answered
- [ ] Lock-in tolerance: how painful switching providers later would be is assessed

C. Cost per learner per month

- [ ] We converted seats, minutes, credits, and tokens to one unit: cost per learner per month
- [ ] We compared that figure to the engineering cost of self-hosting (GPU + ML/infra time)
- [ ] Self-hosting is only chosen past a high, steady-volume break-even — not to 'save fees'

D. Privacy, lock-in, and tracking

- [ ] A GDPR Article 28 data processing agreement is in place; the no-training clause is explicit
- [ ] Sub-processors, retention, and (for US records) FERPA constraints are checked
- [ ] If the feature evaluates learners, EU AI Act high-risk duties are accounted for
- [ ] A thin abstraction layer is in place; the feature emits xAPI statements to our analytics



GO DEEPER · THE 7-VOLUME HANDBOOK

Fora Soft Video Engineering Handbook

Seven volumes, 9,000+ pages: encoding, streaming, audio and AI for video products — from the pixel to production.

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

Browse the handbook on Amazon →

This worksheet walks one AI feature through the four drivers and converts every option to cost per learner per month. Our seven-volume handbook covers the AI side of video products in depth; start with Volume 6 for models, speech and search, then Volume 7 for serving, cost and the EU AI Act, and Volume 3 for delivery.

WHAT YOU WILL FIND INSIDE

- Models, speech recognition and search over video (Vol. 6)
- Players and delivery for course video (Vol. 3)
- Serving, cost per stream-hour and EU AI Act duties (Vol. 7)

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

Pricing an AI tutor, captioning or search feature for your platform?

We will compare product, managed API and self-hosted options for your volumes and give you cost per learner per month for each. 250+ video products shipped since 2005.

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