

Optical Flow Algorithm Picker

Map sparsity, latency, and accuracy requirements to the right primitive — and the right per-frame compute budget.

Decision tree — pick the algorithm

1. Do you need flow at every pixel?

- ☐ **No** (tracking, AR anchors, odometry) → Lucas-Kanade.
- ☐ **Yes** (stabilization, interpolation, VSR) → continue.

2. Is the pipeline real-time?

- ☐ **Yes, CPU only** → DIS (Kroeger 2016, 300–600 Hz).
- ☐ **Yes, GPU available** → SEA-RAFT (20+ FPS at 1080p).
- ☐ **Yes, NVIDIA-only deployment** → NVIDIA Optical Flow SDK.
- ☐ **No, offline** → continue.

3. Accuracy ceiling required?

- ☐ **Standard** → RAFT (Sintel EPE 2.855, KITTI F1 5.10%).
- ☐ **State-of-the-art** → FlowFormer / MegaFlow / DPFlow (2026).

4. Licence requirements?

- ☐ **Commercial OK** → all listed are BSD/Apache/MIT.
- ☐ **NVIDIA hardware lock** → Optical Flow SDK is GPU-locked.

Per-frame compute budget

Algorithm	Type	Hardware	Speed
Lucas-Kanade	Sparse	CPU	~200 FPS @ 720p
Farneback	Dense	CPU	~30 FPS @ 720p
DIS	Dense	CPU	300–600 Hz @ SD
RAFT	Dense	GPU	9 FPS @ 1088×436
SEA-RAFT	Dense	GPU	20+ FPS @ 1080p
FlowFormer	Dense	GPU	~4 FPS @ 1080p

Three-failure-mode audit


- ☐ **Sparsity fits the feature**: dense for stabilization / interpolation / VSR.
- ☐ **Brightness assumption holds**: photometric normalisation under flicker / low light.
- ☐ **Latency budget respected**: 33 ms/frame at 30 FPS — RAFT will not fit.

Library and licence quick reference

- ☐ **Lucas-Kanade / Farneback**: cv2 (Apache 2.0).
- ☐ **DIS**: cv2-contrib DISOpticalFlow (Apache 2.0).
- ☐ **RAFT**: torchvision raft_large (BSD 3-Clause).
- ☐ **SEA-RAFT**: GitHub (BSD 3-Clause).
- ☐ **FlowFormer**: GitHub (MIT).

Companion to: 'Optical Flow Primitives — RAFT vs Lucas-Kanade'

www.forasoft.com/learn/ai-for-video-engineering/articles-ai/optical-flow-raft-vs-lucas-kanade



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 picker maps sparsity, latency and accuracy needs to Lucas-Kanade, DIS, RAFT and newer flow models. Volume 6 covers computer vision for video in depth: tracking, segmentation and pose, model artifacts and licences, and the deployment topologies that decide your per-frame budget.

WHAT YOU WILL FIND INSIDE

- Tracking and motion estimation in video
- CPU, GPU and edge deployment topologies
- Model artifacts and licences for CV models
- Segmentation and pose built on motion

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

Need motion estimation for stabilization, tracking or interpolation?

We will pick the flow method for your latency and hardware, benchmark it on your clips, and integrate it into your pipeline. 250+ video products shipped since 2005.

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