GHOST STREAM

Neural Video Super-Resolution for AV1 Content | A Harvv Feature by AxiomState

harvv.com/speed

+3.92
VMAF vs Lanczos
Real Camera (7 clips)

+4.03
VMAF vs Lanczos
Standard (Median)

47,980
Parameters
(SPAN Architecture)

94 KB
Model Size
(Float16, WebGPU)

Common Questions — Answered With Data

Is this gaming VMAF?
NEG gap +0.18 on camera. PSNR +0.72 dB. Minimal inflation — genuine reconstruction.

Only works on CGI?
12 clips tested (4 CGI + 7 camera + 1 synthetic). Beats Lanczos on all 12.

Only one quality level?
5 CRF points (23-51). Beats Lanczos at every level: +4.01 to +2.83.

Reproducible?
MIT weights + benchmark code public: python benchmark.py --all --download --verify

Standard Test Sequences (Blender Open Movies)

Sequence	Type	Lanczos	Ghost Stream	Delta
Big Buck Bunny	Animation	89.23	90.99	+1.76
Elephants Dream*	CGI / Surreal	72.31	81.33	+9.02
Sintel	CGI / Dark	90.43	93.14	+2.71
Tears of Steel	Live Action + VFX	87.05	92.40	+5.35
Mean (Median)		**84.76**	**89.47**	**+4.71 (+4.03)**

*ED (+9.02) is an outlier; excl. it, 3-clip avg +3.27. Cross-machine 2nd run: +3.81 (SVT-AV1 non-determinism).

Real Camera Footage (Google Pixel Phone)

Clip	Lanczos	Ghost Stream	Delta	NEG Gap
20210119 (outdoor)	87.38	91.10	+3.72	-0.07
20210728 (outdoor)	76.67	80.35	+3.68	+0.13
20220618 (mixed)	78.63	82.94	+4.31	+0.25
20230109 (indoor)	76.26	78.67	+2.41	+0.03
Average (4 clips)	**79.73**	**83.27**	**+3.53**	**+0.09**

7-clip camera avg: +3.92, NEG +0.18. + Mandelbrot synthetic: +3.41. Total: 12 clips benchmarked.

What's Public (github.com/AxiomState/ghoststream-benchmark)

Resource	Contents	License
models/	SPAN weights (94KB float16, ~200KB PyTorch .pth)	MIT
benchmark/	One-command benchmark with --verify flag	MIT
results/	Per-clip VMAF, NEG, PSNR + crop comparisons at 4x	MIT
demo/	WebGPU player, shader, optimizer	MIT
paper/	Technical paper (PDF)	MIT
Training code	Loss functions, distillation, preprocessing	Proprietary

GHOST STREAM

Neural Video Super-Resolution for AV1 Content | A Harvv Feature by AxiomState

harvv.com/speed

Performance Across Quality Levels (BD-Rate)

CRF	Quality	Lanczos	Ghost Stream	Delta
23	Near-lossless	91.56	95.57	+4.01
28	High	89.75	93.69	+3.94
35	Standard	85.98	89.79	+3.81
42	Low	80.87	84.26	+3.39
51	Very low (3G)	73.47	76.30	+2.83

Temporal Consistency

Clip	Ratio	Verdict
Big Buck Bunny	0.999x	Identical
Elephants Dream	1.118x	Acceptable
Sintel	1.015x	Good
Tears of Steel	0.999x	Identical
Pixel 20210119	1.002x	Identical
Pixel 20210728	1.012x	Good

Ablation

Configuration	vs Lanczos	Gain
Charbonnier only	-2.00	baseline
+ Adversarial baseline	-0.98	+1.02
+ Distillation (generic)	-0.34	+0.64
+ Diverse data (DIV2K)	+3.92	+4.26

Rows 1-3: EU-CZ-1. Row 4: US-NC-1. Cross-machine variance +/-0.9.

Independent Verification

> PSNR improves +0.72 dB alongside VMAF on every clip — proves genuine reconstruction, not hallucinated detail
> VMAF NEG gap +0.18 on camera (negligible inflation). CGI gap +0.78 (mild enhancement, disclosed)
> Reproducibility: python benchmark.py --all --download --verify (tolerance +/-1.0 for cross-machine SVT-AV1 variance)
> Result reproduced across 2 training runs, different seeds: +/-0.07 VMAF. Cross-machine delta variance: +/-0.9
> Model weights MIT licensed. Training code proprietary. Benchmark code + all results public.

Estimated Impact

Traditional 720p: ~2 Mbps. Ghost Stream 360p + client SR: ~250 kbps. 6-8x bandwidth reduction.

For 1M monthly viewers x 2hr: est. $378K/year CDN savings.

Methodology & Transparency

All results: Netflix VMAF + VMAF NEG. SVT-AV1 non-deterministic across machines; deltas consistent +/-0.9.
Total R&D: ~$200 GPU compute ($120 was idle pods before auto-shutdown). Per-run training: $0.50.

Free Tool: harvv.com/speed
API Docs: docs.harvv.com
Benchmarks: github.com/AxiomState/ghoststream-benchmark
Contact: jordan@harvv.com

AxiomState builds tools that make e-commerce faster and smarter.
Harvv detects friction. Ghost Stream eliminates video friction. Zero PII.