

8 WEEKS → TIMELINE COMPRESSED

Compressed Feasibility

A technical update to the ambient non-consensual synthesis threat model.

Eight weeks did not change the threat. They changed its clock.

Three independently released open-source tools collapsed the feasibility timeline we published in January. The edge-only class we projected at 12 months may now arrive in **6–9. The window for proactive policy has narrowed accordingly.**

None of the three teams built their tool for this purpose. None appear aware of the pipeline they collectively enable.

SOURCES Gilly (2026), Ambient Non-Consensual Image Synthesis, Zenodo, doi:10.5281/zenodo.18297286.

The January model: two classes feasible, two on the horizon.

CLASS 1 • NOW Cloud-assisted real-time Feasible at publication.	CLASS 2 • NOW Edge static + interpolation Feasible at publication.	CLASS 3 • +12 MO Edge full video synthesis Projected, moderate assumptions.	CLASS 4 • +24-36 MO Edge real-time <500ms Projected, continuous synthesis.
--	---	--	--

The harm is borne overwhelmingly by women. This update revisits the timeline, not the nature of the threat, which the original analysis establishes in full.

Three components, each solving a legitimate problem,
composable into one.

STAGE • SEGMENTATION

MatAnyone2

140 MB

SOTA video matting, clean masks on wearable-class hardware.

STAGE • INFERENCE

Diagonal Distillation

270×

Speedup over baseline, 5s of video in 2.6s.

STAGE • RENDERING

MobileGS

4.8 MB

120+ FPS Gaussian splatting on a Snapdragon phone.

The bottleneck that dominated the model is exactly where the time collapses.


PIPELINE STAGE	ORIGINAL	REVISED	BASIS
Person detection + segmentation	35–60ms	15–30ms	MatAnyone2 quality at speed
Region crop + pose estimation	20–30ms	15–25ms	Cleaner masks, less post-proc.
Synthesis model inference	100–500ms	40–200ms	DiagDistill trajectory
Video interpolation / consistency	100ms–5s	50ms–2s	Long-sequence stability

SOURCES Camera input (<1ms) and AR display composition (5–10ms) unchanged. Yang et al., CVPR 2026 (arXiv:2512.11782); Liu et al., ICLR 2026 (arXiv:2603.09488).

The projected classes pull months closer.

Edge full video synthesis

DiagDistill exceeds projected speedup



Edge real-time <500ms

MobileGS + DiagDistill trajectory



The direction of compression is unambiguous.

SOURCES Liu et al., ICLR 2026 (arXiv:2603.09488); Du et al., ICLR 2026 (arXiv:2603.11531).

Yesterday's "breakthrough" case is today's baseline.

The observed trajectory matched our original **aggressive** scenario, so it becomes the revised **moderate** one. A trajectory projection, not a guarantee: thermal limits and sustained-vs-benchmark gaps remain. But the shift happened in eight weeks.

NEW AGGRESSIVE SCENARIO

Edge full video synthesis **3-6 mo**

Edge real-time <500ms **12-18 mo**

Capability matures in weeks. Governance answers in years.

Capability

6–9 months to edge feasibility

Legislation

3–5 years to enact

Every month of delay brings the edge-only class, **no network trace, no external service, fully deniable**, closer to feasibility, after which enforcement becomes substantially harder.

The components exist. They are free. They are documented.

The only remaining barrier is the glue code that connects them into a pipeline, **and that is precisely the task at which current large language models excel**. The last obstacle is the one AI is now best at removing.


open source

deployment-ready code

documented

hosted demos

No component is harmful. The composition is.



Platform review that evaluates models only in isolation is structurally blind to combinatorial threats.

Three moves the compressed timeline demands.

01

Legislate creation, not distribution

Edge-only deployment leaves no trace, distribution-stage enforcement arrives too late.

02

Govern pipeline composition

Hosting platforms must evaluate combinations, not just individual models.

03

Monitor continuously

Eight-week shelf life: static risk assessments will perpetually lag the capability.

SOURCES The original recommendations remain applicable. These respond specifically to the compression. Gilly (2026), Compressed Feasibility, sec. 5.

No one cited the threat. No one designed the pipeline. It assembled itself anyway.

Eight weeks produced a **270×** speedup, a **140 MB** segmenter, and a **4.8 MB** renderer, from teams unaware of one another. This is exactly the pattern our analysis predicted, and exactly the pattern current governance is not built to detect.

The tools change the timing. The governance gap does the rest.

SOURCES Gilly, T. (2026). Compressed Feasibility: A Technical Update to the Ambient Non-Consensual Synthesis Threat Model. Real Safety AI Foundation.

— REFERENCES

Gilly, T. (2026). *Compressed Feasibility: A Technical Update to the Ambient Non-Consensual Synthesis Threat Model*. Real Safety AI Foundation.

[1] T. Gilly, "Ambient Non-Consensual Image Synthesis: A Convergence Threat Analysis of AR Wearables, Real-Time Video Generation, and Open-Source Nudification Models," Zenodo, Jan. 2026.
<https://doi.org/10.5281/zenodo.18297286>

[2] P. Yang, S. Zhou, K. Hao, and Q. Tao, "MatAnyone 2: Scaling Video Matting via a Learned Quality Evaluator," in Proc. CVPR, 2026. arXiv:2512.11782. GitHub: <https://github.com/pq-yang/MatAnyone2>. HuggingFaceDemo:<https://huggingface.co/spaces/PeiqingYang/MatAnyone2>.

[3] J. Liu, X. Liu, K. Mei, Y. Wen, M.-H. Yang, and W. Liu, "Streaming Autoregressive Video Generation via Diagonal Distillation," in Proc. ICLR, 2026. arXiv:2603.09488. GitHub: <https://github.com/Sphere-AI-Lab/diagdistill>. Project page:<https://spherelab.ai/diagdistill/>.

[4] X. Du et al., "Mobile-GS: Real-time Gaussian Splatting for Mobile Devices," in Proc. ICLR, 2026. arXiv:2603.11531. Project page: <https://xiaobiaodu.github.io/mobile-gs-project/>.

[5] InSpatio-WorldFM: An Open-Source Real-Time Generative Frame Model for Spatial Intelligence. arXiv:2603.11911, March 2026.

[6] T. Gilly, "Emergent Convergence Risk: Why Existing AI Governance Cannot Detect Combinatorial Threats and a Proposed Methodology," forthcoming, 2026. Word Count: Approximately 2,200 words (excluding references)