

Ambient Non-Consensual Image Synthesis

A convergence threat analysis of AR wearables, real-time video generation, and open-source nudification models.

Four domains. Studied separately. Converging now.

Open-source nudification models, real-time video generation, consumer AR wearables with integrated cameras, and edge AI inference have each attracted scholarly attention in isolation. No prior work examines what happens when they converge: when someone wearing commercially available glasses can, in real time, render synthetic intimate imagery of fully clothed people in their field of view.

JANUARY 2026 · GROK INCIDENT

Within days, coordinated users produced hundreds of thousands of non-consensual intimate images of real people, including minors. The California Attorney General launched a formal investigation.

The demand is not hypothetical. The question is whether the tools will remain platform-mediated, or become ambient.

Each stage contracts the domain of safe participation.

STAGE 1 · PRE-DIGITAL Physical access required Non-consensual sexualization required physical proximity or physical photographs. All unobserved spaces were safe.	STAGE 2 · EARLY INTERNET Voluntary disclosure bounded risk Risk bounded by what was shared. Physical spaces and unshared digital content remained safe.	STAGE 3 · DEEPPFAKE ERA (NOW) Any shared image is source material Documented chilling effects: women withdraw from online participation to avoid targeting. 102 undressing requests per 10 minutes in the Grok incident.	STAGE 4 · AMBIENT AR (IMMINENT) Physical presence is sufficient No photograph needed. Existing in view of an AR-equipped observer provides source material for real-time synthesis. Safe public participation is eliminated.
---	---	--	--

The Convergence

Four mature components. Cloud-assisted: feasible today. On-device: twelve to twenty-four months.



Real-time ambient non-consensual intimate image synthesis is architecturally feasible today via cloud-assisted pipelines and will be feasible on-device within 12 to 24 months given current hardware and model trajectories.

All four components are commercially available.

Open-source nudification models

More than twenty commercial applications; thousands of models on public hosting platforms. Per-image cost: zero to six cents. No technical barriers to entry.

Consumer AR wearables

Meta Ray-Ban glasses: 12-megapixel cameras, microphones, AI assistant in conventional eyewear form. Capable of continuous video capture. LED indicator can be obscured.

Real-time video generation

PixVerse R1 (Jan. 2026): interactive 1080p streaming, 1 to 4 diffusion steps. TurboDiffusion: 100 to 200 times speedup. SnapGen-V: 5-second video in 5 seconds on mobile.

Edge and cloud AI inference

Real-time computer vision on mobile hardware already achieved. Cloud-assisted pipeline feasible today; on-device latency targets will be met within 12 to 24 months.

Perceptual consent: autonomy over how one's body is rendered in another's perception.

Traditional consent frameworks require an act of disclosure: sharing a photograph, publishing a record. Ambient synthesis violates bodily autonomy without any act by the subject. The harm is existence in view. The concept is grounded in Nussbaum's capabilities approach, Solove's privacy taxonomy, and Kantian human dignity.

THREE DISTINGUISHING CRITERIA

Realism

The rendering is realistic enough to function as a representation of the subject's body.

Externalization

It exists in rendered form outside the mind, in a visual field or a file.

Specificity

It targets an identifiable individual, not an imagined or composite person.

Mitigations Fail

No technical mitigation we have examined is sufficient to prevent ambient synthesis by motivated attackers.



On-device safety filters are fundamentally vulnerable to adversarial bypass. Once bypassed by one skilled attacker, the exploit becomes a distributable tool requiring no technical skill to deploy.

Defenders must succeed always. Attackers need succeed once.

Neural networks are inherently susceptible to adversarial inputs: modifications imperceptibly different from normal inputs that produce arbitrarily different outputs. Any filter on consumer hardware can be extracted through reverse engineering, attacked, and distributed. The defenses are not independent: an attacker who moves to local processing automatically bypasses platform moderation, and the exploit requires no technical skill from subsequent users.

SAFETYCORE · A DEPLOYED CASE

Successful attacks against Google's SafetyCore on-device safety system have been publicly documented and described in peer-reviewed preprints.

The asymmetry is not one attacker versus one defender. It is one skilled attacker enabling millions of unskilled users. Friction-based defenses collapse entirely once bypass tools circulate.

Regulatory Gaps

Every existing law targets distribution. Perception-only synthesis is unaddressed.



The EU AI Act, US federal statutes, and UK law all share the same structural gap: they require distribution as an element of the violation. A synthesis that never leaves the perpetrator's visual field is not reached.

Every major framework requires distribution. Ambient synthesis does not distribute.

EU AI Act (2024/1689)	Biometric provisions address database creation, not real-time perceptual manipulation. No provision requires disclosure of AI model capabilities on consumer AR hardware.
DEFIANCE Act (US)	Civil cause of action for victims. Requires distribution as the harm vector. Creation without distribution is not covered. Enforcement requires the victim to become aware.
TAKE IT DOWN Act (US)	Criminal penalties and platform takedown obligations. Also assumes distribution as the harm vector. Perception-only synthesis not addressed.
UK Online Safety Act 2023	Section 66B explicitly covers AI-generated imagery but requires "sharing." Voyeurism statutes address equipment "beneath clothing," not synthesis from clothed imagery.

SOURCES Regulation (EU) 2024/1689 [30]. · DEFIANCE Act, H.R. 3562 [31]. · TAKE IT DOWN Act (2025) [32]. · Online Safety Act 2023, s.66B [33]. · Voyeurism (Offences) Act 2019 [59].

The population-level harm does not require individual victimization.

Empirical research documents chilling effects on women's participation among women who have never been individually targeted. Anticipatory threat and collective cultural injury operate at the population scale, independent of whether any individual has been harmed. Stage 4 extends these documented harms from digital to physical space, and from asynchronous capture to real-time perception.

ICFJ, THE CHILLING (2023)

Systematic self-censorship and withdrawal from professional online activity by women journalists, documented globally.

CHAPMAN & WILLIAMS (2025)

Normalization documented among users of nudification platforms, suggesting ambient deployment would face similar social rationalization pressures.

Recommendations

Technical, legal, and research responses that match the actual threat.

IV

The timeline mismatch is real: cloud-assisted synthesis is feasible now; legislative processes require three to five years. Industry self-regulation, professional norms, and civil society must fill the gap.

Three reforms, none technical alone.

LAW

Criminalize synthesis without distribution.

Extend SNEACI statutes to cover rendering and generating intimate imagery of an identifiable person without consent, regardless of whether the result is stored, distributed, or shared. Require mandatory capability disclosure from AR hardware manufacturers.

TECHNICAL

Cross-domain threat modeling.

Convene SNEACI researchers, XR privacy scholars, and video generation experts for joint threat analysis. Prioritize adversarial robustness for wearable safety systems. Develop model extraction detection and prevention for wearable platforms.

RESEARCH

Formal theory of perceptual consent.

Ground the concept in XR privacy scholarship and SNEACI harm research. Develop hardware and OS-level detection mechanisms. Conduct empirical prevalence studies on psychological impacts of perceptual consent violations.

The question is whether frameworks will be in place before harm accumulates.

The threat is not speculative. The January 2026 Grok incident demonstrated immediate scaled demand when tools become accessible. The component technologies are publicly documented, commercially available, and on clear development trajectories. The question is not whether ambient AR-based synthesis will be attempted. It is whether the technical, legal, and research communities will have coordinated a response before harm normalizes. The pattern of reactive response after capability has matured and harm has accumulated need not repeat.

— REFERENCES

Gilly, T. (2026). *Ambient Non-Consensual Image Synthesis: A Convergence Threat Analysis of AR Wearables, Real-Time Video Generation, and Open-Source Nudification Models*. Real Safety AI Foundation.

[1] C. Gibson et al., "Analyzing the AI Nudification Application Ecosystem," in Proc. USENIX Security Symposium, 2025.

[2] PixVerse, "PixVerse R1: Next-Generation Real-Time World Model," Technical Report, Jan. 2026. [Online]. Available: <https://pixverse.ai/en/blog/pixverse-r1-nextgeneration-real-time-world-model>

[3] "PixVerse R1: Real-Time AI Video Generation Arrives," VP Land, Jan. 2026.

[4] "AIsphere Unveils PixVerse R1, Ushering in Real-Time AI Video Era," TMTPost, Jan. 2026.

[5] "Smart glasses: 'I was secretly filmed and trolled online,'" BBC News, Jan. 2026.

[6] "Meta's AI-Powered Smart Glasses Raises Privacy & User Data Concerns," Carleton University, 2025.

[7] "The Role of Edge AI in Real-Time Analytics in 2026 Explained," Kanerika, 2026.

[8] "How Edge AI Enables Real-Time Video Processing in Smart Cameras," TechNexion, 2025.

[9] "Virtual Reality Data and Its Privacy Regulatory Challenges: A Call to Move Beyond Text-Based Informed Consent," California Law Review, 2023.

[10] "Speculative Privacy Concerns About AR Glasses Data Collection," Proc. Privacy Enhancing Technologies (POPETS), 2023.

[11] "SoK: Data Privacy in Virtual Reality," arXiv:2301.05940, 2023.

[12] ShengShu Technology and Tsinghua University, "TurboDiffusion: Accelerating Video Diffusion Models by 100-200 Times," arXiv:2512.16093, Dec. 2025. [Online]. Available: <https://arxiv.org/abs/2512.16093>

[13] "SnapGen-V: Generating a Five-Second Video within Five Seconds on a Mobile Device," arXiv:2412.10494, 2024.

[14] "When On-Device AI Becomes a Security Flaw: The SafetyCore Case," Skyld AI, 2025.

[15] "Breaking SafetyCore: Exploring the Risks of On-Device AI Deployment," arXiv:2509.06371, 2025.

— REFERENCES (CONTINUED)

[16] M. C. Nussbaum, Women and Human Development: The Capabilities Approach. Cambridge University Press, 2000.

[17] D. J. Solove, "A Taxonomy of Privacy," University of Pennsylvania Law Review, vol. 154, no. 3, pp. 477-564, 2006.

[18] I. Kant, Groundwork of the Metaphysics of Morals, 1785. M. Gregor, Trans., Cambridge University Press, 1998.

[19] "Unmasking digital deceptions: An integrative review of deepfake detection," PMC, 2025.

[20] "'There Are No Limits!': AI Undressing Apps and the Normalization of Non- Consensual Intimate Imagery," PubMed, 2025.

[21] "A systematic literature review on Virtual Reality and Augmented Reality in terms of privacy, authorization and data-leaks," arXiv:2212.04621, 2022.

[22] "Exploring the Uncoordinated Privacy Protections of Eye Tracking and VR Motion Data," arXiv:2411.12766, 2024.

[23] "Behind Grok's mass digital undressing lies an 'unsurprising' cohort," ABC News Australia, Jan. 2026.

[24] "Attorney General Bonta Launches Investigation into xAI, Grok," California Attorney General Press Release, Jan. 2026.

[25] "Grok-created images of real people in bikinis, underwear banned on X," Mashable, Jan. 2026.

[26] A. Shostack, Threat Modeling: Designing for Security. Wiley, 2014.

[27] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and Harnessing Adversarial Examples," in Proc. ICLR, 2015. arXiv:1412.6572.

[28] N. Carlini and D. Wagner, "Towards Evaluating the Robustness of Neural Networks," in Proc. IEEE S&P, 2017.

[29] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, and A. Swami, "The Limitations of Deep Learning in Adversarial Settings," in Proc. IEEE Euro S&P, 2016.

[30] European Parliament and Council, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act)," Official Journal of the European Union, Aug. 2024. EUR-Lex: 32024R1689.

[31] "DEFIANCE Act of 2025," H.R. 3562, 119th Congress (2025-2026). [Online]. Available: <https://www.congress.gov/bill/119th-congress/house-bill/3562>

[32] "TAKE IT DOWN Act," Pub. L. No. 119-XX, May 2025.

[33] "Online Safety Act 2023," UK Public General Acts, c. 50, Section 66B.

[34] N. Henry et al., "Image-Based Sexual Abuse: A Study on the Causes and Consequences of Non-Consensual Nude or Sexual Imagery," 2020.

[35] "Prevalence and Impacts of Image-Based Sexual Abuse," Cyber Civil Rights Initiative, 2025.

— REFERENCES (CONTINUED)

[36] "Development of an AI-powered AR glasses system for real-time first aid assistance," PMC, 2025.

[37] E. Maiberg, "Hugging Face Is Hosting 5,000 Nonconsensual AI Models of Real People," 404 Media, July 15, 2025. [Online]. Available: <https://www.404media.co/hugging-face-is-hosting-5-000-nonconsensual-aimodels-of-real-people/>

[38] "UnfilteredAI," Hugging Face Model Repository.

[39] "Researchers warn of rise in AI-created non-consensual explicit images," University of Florida News, May 2025.

[40] "The 10 Best AI Video Generators for Content Creation," Digital Bricks, 2025.

[41] "The Best AI Video Creation Trends from 2025," Clippie AI, 2026.

[42] "PixVerse V5.5: Multi-Shot Video with Sound," SuperMaker AI, 2025.

[43] "PixVerse-R1 LLM Benchmark Data," Automatio AI, 2026.

[44] "Meta AI Glasses: Revolutionary AR Technology or Privacy Nightmare?" GigeNet, 2025.

[45] M. Speicher et al., "AR Privacy Concerns in Social Contexts," in Proc. ACM CHI, 2020.

[46] R. Williams et al., "Bystander Privacy in AR," IEEE VR, 2021.

[47] "Edge Computing and AI: The Future of Real-Time Data Processing," Sapien.io, 2025.

[48] "AI-based Wearable Vision Assistance System for the Visually Impaired," arXiv:2412.20059, 2024.

[49] "LoXR: Performance Evaluation of Locally Executing LLMs on XR Devices," arXiv:2502.15761, 2025.

[50] "Immersed in Reality Secured by Design: A Comprehensive Analysis of Security Measures in AR/VR Environments," arXiv:2404.16839, 2024.

[51] "Deepfakes and Synthetic Media," TechUK, 2025.

[52] "The rise and risks of synthetic media," Digital Watch Observatory, 2025.

[53] "2025 AI Safety Index," Future of Life Institute, 2025.

[54] "SDXS: Real-Time One-Step Latent Diffusion Models with Image Conditions," arXiv:2403.16627, 2024.

[55] "'Violation of my body:' Perceptions of AI-generated non-consensual (intimate) imagery," arXiv:2406.05520, 2024.

[56] "Regulating AI Deepfakes and Synthetic Media in the Political Arena," Brennan Center for Justice, 2025.

[57] J. Bailenson and J. Cummings, "VR Consent and Embodiment," Stanford Virtual Human Interaction Lab, 2016.

[58] "Designing Effective Consent Mechanisms for Spontaneous AR Interactions," in Proc. ACM CHI, 2025.

— REFERENCES (CONTINUED)

[59] "Voyeurism (Offences) Act 2019," UK Public General Acts, c. 2, Section 67A.

[60] California Penal Code § 647(j)(4).

[61] California Civil Code § 1708.85.

[62] Virginia Code § 18.2-386.2.

[63] Texas Penal Code § 21.165.

[64] "Google forcing Android System SafetyCore on users to scan for nudes," Kaspersky, 2025.

[65] "Smart Glasses: Cool Tech or a Privacy Threat?" Panda Security, 2025.

[66] "General Data Protection Regulation," Regulation (EU) 2016/679, Articles 4(14), 9.

[67] Council of Europe, "Convention on preventing and combating violence against women and domestic violence (Istanbul Convention)," CETS No. 210, 2011.

[68] "Not a Good Look, AI: What Happens to Privacy When Glasses Get Smart," Cyber Security Advisors Network, May 2025.

[69] "Owl-1: Omni World Model for Consistent Long Video Generation," arXiv:2412.09600, 2024.

[70] "Augmenting Security and Privacy in the Virtual Realm: An Analysis of Extended Reality Devices," arXiv:2402.03114, 2024.

[71] International Center for Journalists, "The Chilling: A Global Study of Online Violence Against Women Journalists," 2023. [Online]. Available: <https://www.icfj.org/our-work/chilling-global-study-online-violence-againstwomen-journalists>

[72] M. Björkenfeldt et al., "Self-Censorship Among Journalists: Anticipatory Threat and Professional Adaptation," Digital Journalism, 2025. DOI: 10.1080/21670811.2025.2601961

[73] WITNESS, "Technology-Facilitated Gender-Based Violence," Mar. 2025. [Online]. Available: <https://blog.witness.org/2025/03/technology-facilitated-genderbased-violence/>

[74] UNFPA, "An Infographic Guide to Technology-Facilitated Gender-Based Violence," 2025. [Online]. Available: <https://www.unfpa.org/>

[75] L. Lumsden et al., "Gender and Online Expression: A UK Population Study," arXiv:2403.19037, 2024.

[76] Washington Post, "X Users Tell Grok to 'Undress' Women in Their Photos," Jan. 6, 2026.

[77] The Conversation, "How AI-Generated Sexual Images Cause Real Harm, Even Though We Know They Are Fake," Jan. 15, 2026. [Online]. Available: <https://theconversation.com/how-ai-generated-sexual-images-cause-real-harmeven-though-we-know-they-are-fake-273427>

[78] F. Martin, "Online Safety Regulation of Deepfake and Image-Based Sexual Abuse in Australia," Alternative Law Journal, 2025. DOI: 10.1080/10383441.2025.2504791

— REFERENCES (CONTINUED)

[79] A. Chapman and M. Williams, "'There Are No Limits!': AI Undressing Apps and the Normalization of Nonconsensual Intimate Deepfakes," Violence Against Women, 2025. DOI: 10.1177/10778012251397966

[80] A. Singh, "The Gendered Panopticon: How Surveillance Enforces Binary Norms," Nickled and Dimed, Apr. 2025. [Online]. Available: <https://nickledanddimed.com/2025/04/21/the-gendered-panopticon-howsurveillance-enforces-binary-norms/>

[81] "The Panopticon 2.0: Mass Surveillance in the Digital Age," The Uncensored Truth, 2025. [Online]. Available: <https://theuncensoredtruth.com/part-1-thepanopticon-2-0-mass-surveillance-in-the-digital-age/>

[82] M. Vef and B. Groß, "NUCA: AI Camera That Redefines Image Creation," 2024. [Online]. Available: <https://nuca.rocks/Correspondingauthor:TravisGilly,RealSafetyAIFoundation,t.gilly@ai-literacy-labs.org>