The wavefield gate: Fourier modes are the wrong basis for visual computing

W. Schulz

Sep 8, 2026

Abstract

A proposed photonic 'wavefield' graphics architecture, tested on its two measurable claims. Fourier modes need 21–49% of coefficients to reach 35 dB, not 0.5%, and a one-pixel move changes 97% of them. Display space–bandwidth and conversion costs close the rest: the physics is sound, the advantages are not.

On 8 September 2026 a conversation thread arrived that began with a familiar animation — rotating circles, each riding on the last, tracing a heart — and escalated over seven rounds into a proposal for an end-to-end visual-computing architecture: light to modes to bits to modes to light, with no framebuffer, a "wavefield compiler", a modal codec that sends only coefficient changes, and photonic hardware to do the work. The project ran it through both of its gates the same day: the novelty gate, scored on both axes as the soundness-axis amendment requires (a claim can be old and true, and "known" is to be reported as "known and therefore probably true", not as a deflation), and the substrate gate, which asks whether a hardware story survives division.

The verdict was CLOSED, no build. The physics is sound — it is Fourier optics and computer-generated holography, a live field. Every stated *advantage* is either owned, inverted, or forbidden by a conservation law. Nothing in the thread was younger than the literature.

Every layer has an owner

Proposed	Owner
Epicycles draw any closed contour	Fourier descriptors of closed curves — Zahn & Roskies, 1972 [1]; elliptic Fourier features — Kuhl & Giardina, 1982 [2]
"Program the waves, not the pixels"	The definition of a hologram; computer-generated holograms — Lohmann & Paris, 1967 [3]
A wavefield compiler	Iterative phase retrieval — Gerchberg & Saxton, 1972 [4] — and its gradient-descent successors
A learned, graphics-native optical basis	Neural holography — Peng, Choi, Padmanaban & Wetzstein, 2020 [5]

Proposed	Owner
Foveated wavefield rendering with eye tracking	Holographic near-eye displays with eye-tracked approximations — Maimone, Georgiou & Kollin, 2017 [6]; foveated holographic displays — Chakravarthula et al., 2021 [7]
A sensor that measures coefficients directly	Compressive sensing — Candès, Romberg & Tao, 2006 [8]; Donoho, 2006 [9]; the single-pixel camera — Duarte et al., 2008 [10]
A modal codec that sends coefficient deltas	Transform coding — the discrete cosine transform, Ahmed, Natarajan & Rao, 1974 [11], and every codec since
The scene as compact, resolution-independent math	Vector graphics, from Sketchpad, 1963 [12]; a triangle mesh already is this
Photonic matrix–vector hardware for the coefficients	Already closed in this project's photonics track, where analog in-memory compute beat it by roughly 10.8× on latency and 20.9× on energy

The two falsifiable claims, measured

Two of the thread's claims could be tested directly, and they were, in a forty-line harness pre-registered in its own header before it ran: four orthonormal bases (pixel, DFT, DCT, Haar) on 512×512 targets, oracle top-K coefficients, so every budget is exactly comparable.

"A few thousand modes can replace millions of pixels." The fraction of coefficients each basis needs to reach 35 dB PSNR, in this project's run:

Target	Pixel	DFT	DCT	Haar
The demo's own heart, as a 2-D field	1.24%	49.0%	19.5%	1.55%
Text	9.4%	24.6%	19.8%	2.16%
A synthetic natural-image stand-in	96.9%	21.6%	15.3%	13.5%

At the pitch's own ratio — ten thousand of two million, 0.5% — the best any basis reached on any target was 28.7 dB. Fourier modes miss the pitch by 40 to 100 times and were the worst transform on every target. The localised bases win on edges — and localised bases are exactly what a plane-wave optical mode does not implement. One registered prediction missed and is reported as such: at the smallest budget, 0.1%, the localised bases did *not* beat Fourier on the edge targets, because 262 isolated pixels or wavelets lose to a Fourier blur. The natural-image target is synthetic (a $1/f^2$ field with discs); the repository holds no photograph, and that is a corpus caveat on the third row.

"When an object moves, only a few coefficients change." A disc of radius 12, shifted by one pixel:

Basis	Coefficients changed	Update energy $\|\Delta c\|^2/\|c\|^2$
Pixel	0.18%	0.1408
Haar	0.08%	0.1408
DCT	79.1%	0.1408
DFT	97.2%	0.1408

The sign is backwards. By the shift theorem a local edit multiplies every Fourier coefficient by a phase ramp, so Fourier maximally delocalises it. And the update energy is identical in every basis, as Parseval's theorem requires; the claim could only ever have been about sparsity, never about work.

The demo looks miraculous for a precise reason. Eight circles reproduce the *one-dimensional* parametric heart curve to a relative error of about 3×10^{-16}. The thread then silently swapped that 1-D curve, which is compact in Fourier, for the 2-D image of the same heart, which needs 49% of its coefficients. That swap is the whole illusion.

The arithmetic

Space–bandwidth product. The number of pixels a holographic display must control is fixed by eyebox, field of view and wavelength, independent of content and of representation [13]. At 532 nm and a 60° field, an eye-tracked 4 mm eyebox needs about 5.7×10^7 pixels; an untracked 10 mm eyebox needs 3.5×10^8; a 10 cm tabletop needs 3.5×10^{10}. The best phase-only modulator the gate priced, a 10-megapixel panel at a 3.74 μm pitch [14], supports a field of about 8.2° and falls short of those three cases by 5.5×, 34× and 3447×. Foveation lowers the number of *content* bits, not the number of modulator pixels. The modulator frame for the eye-tracked case moves 41 Gbit/s per eye — it *is* a framebuffer, of complex numbers.

Where the compute goes. Per eye-frame at that size, one FFT propagation is about 7×10^9 floating-point operations; gradient-descent hologram synthesis at 200 iterations is about 3×10^{12}, or 262 TFLOP/s at 90 Hz. The multiplier is the phase-only constraint on the modulator, which turns synthesis into non-convex retrieval. That is a device boundary. Basis choice does not touch it.

The boundary tax. "Let the lens do the Fourier transform" replaces a digital 2-D FFT of about $2.5{\cdot}N{\cdot}\log_2 N$ multiply-accumulates with N optical readouts: 52–62 MACs per conversion at 2–33 megapixels, below the substrate gate's breakeven of 111 for an 8-bit conversion. And the lens's Fourier-transform property is a property of *coherent* optics [15]; natural scenes are incoherent, so the pixels must be captured first. The camera-side idea dies twice.

Six more rounds, and one correction

The thread did not stop, and each round was answered on the same terms.

- A photonic state machine — visual state held in light, updated by sparse commands. Light does not sit still: even a ring resonator with Q of 10^9 holds a photon for under a microsecond, against an 8.3 ms frame. What persists is the modulator's drive state, which is electronic. A K-channel controller feeding any fixed linear optical transform reaches only a K-dimensional subspace of the display's field space, so "ten thousand modes drive a million elements" fails by rank. Whole-field motion can be done passively with prisms and mirrors — that is optical timewarp, and holographic displays already shift their exit pupil to follow a tracked eye [16] — but moving one object

while another stays, or occluding anything, requires acting differently on overlapping parts of one superposition, which linear optics cannot do.

- Which graphics operations beat the boundary? Priced individually — 2-D FFT, separable and non-separable convolution, whole-field warps — none clears 111 MACs per conversion mid-pipeline. The only zero-boundary optical stage is the passive optics *after* the modulator: the lens and the prism, already in every holographic display.
- Map Gaussian Wave Splatting onto photonics. Here the gate corrected itself, visibly. Round 3 had priced single transforms, not hologram synthesis as a whole, which is far denser: the fast Gaussian Wave Splatting method [17] is roughly G = 15,000 MACs per pixel at 15,000 Gaussians, and neural synthesis about 500. A photonic synthesis stage therefore sits at 250–7,500 MACs per conversion, and the boundary-tax check passes. What kills the architecture instead is double-counting: the propagation optics performs "for free" is free only if each primitive physically sits at its depth. Either depth is encoded as phase and computed electronically — the work the stage was meant to remove — or the primitives are written to a stack of depth planes, which is the multifocal display [18][19][20], whose known unsolved defect is occlusion between planes.
- Would a wave basis make learning more efficient? For a linear first layer, $W\cdot(Ux) = (WU)\cdot x$: a fixed orthonormal basis is absorbed into the weights. It can act only through truncation, inductive bias, or optimiser conditioning, and each is owned — Fourier neural operators win where the operator is Fourier-diagonal [21]; spectral bias and Fourier features are preconditioning [22][23]; learning directly from DCT-domain JPEG data was about speed, not a learning advantage [24][25]. Not run, because the identity and exp01 already decide it.
- Where does photonics genuinely win in holographic graphics? Ranked: the display itself, passive steering optics, and the link to the headset. Compute comes last. A computed hologram is a modulator drive pattern — numbers, not light — so a photonic hologram accelerator must convert its output back to numbers, exactly the boundary it claimed to avoid.
- A holographic projector? Real and shipping — the Envisics holographic head-up display in Jaguar Land Rover vehicles is the clearest mass-market case [26] — and its advantages are optical efficiency and focus, not compute. On an ordinary diffuse wall, phase is destroyed and the wavefield advantages go with it.

A wording amendment closed the gate. An earlier line in the session — that nobody works on "which basis" any more — was too absolute and was withdrawn: scene representation for holography is an active research area. What the gate kills is narrower: that an arbitrary visual field is a small number of Fourier modes, and that visual state can persist in light under sparse commands.

What survives

On the soundness axis, a great deal: Fourier optics, computer-generated holography and holographic near-eye displays are sound, corroborated and live. Their open problems — modulator space–bandwidth product, the phase-only constraint, speckle, the eyebox–field-of-view trade — are device problems, none of them representational, and none of them this project's. The gate leaves behind one reusable instrument, the equal-budget basis harness, for the next new-basis pitch, and one recorded lesson: the thread escalated across seven rounds without computing a number. Two divisions — the space–bandwidth product and MACs per conversion — and one distinction between a 1-D curve and a 2-D field settled it.

Decision: close; do not write the specification.

Sources

1. C. T. Zahn and R. Z. Roskies, “Fourier Descriptors for Plane Closed Curves,” *IEEE Transactions on Computers* C-21(3):269–281, 1972. DOI: 10.1109/TC.1972.5008949
2. F. P. Kuhl and C. R. Giardina, “Elliptic Fourier features of a closed contour,” *Computer Graphics and Image Processing* 18(3):236–258, 1982. DOI: 10.1016/0146-664X(82)90034-X
3. A. W. Lohmann and D. P. Paris, “Binary Fraunhofer Holograms, Generated by Computer,” *Applied Optics* 6(10):1739, 1967. DOI: 10.1364/AO.6.001739
4. R. W. Gerchberg and W. O. Saxton, “A practical algorithm for the determination of phase from image and diffraction plane pictures,” *Optik* 35:237–246, 1972. https://cir.nii.ac.jp/crid/1570854175282689408
5. Y. Peng, S. Choi, N. Padmanaban and G. Wetzstein, “Neural holography with camera-in-the-loop training,” *ACM Transactions on Graphics* 39(6), art. 185 (SIGGRAPH Asia), 2020. DOI: 10.1145/3414685.3417802
6. A. Maimone, A. Georgiou and J. S. Kollin, “Holographic near-eye displays for virtual and augmented reality,” *ACM Transactions on Graphics* 36(4) (SIGGRAPH), 2017. DOI: 10.1145/3072959.3073624
7. P. Chakravarthula, Z. Zhang, O. Tursun, P. Didyk, Q. Sun and H. Fuchs, “Gaze-Contingent Retinal Speckle Suppression for Perceptually-Matched Foveated Holographic Displays,” *IEEE Transactions on Visualization and Computer Graphics* 27(11):4194–4203, 2021. DOI: 10.1109/TVCG.2021.3106433
8. E. J. Candès, J. Romberg and T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory* 52(2):489–509, 2006. DOI: 10.1109/TIT.2005.862083
9. D. L. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory* 52(4):1289–1306, 2006. DOI: 10.1109/TIT.2006.871582
10. M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly and R. G. Baraniuk, “Single-pixel imaging via compressive sampling,” *IEEE Signal Processing Magazine* 25(2):83–91, 2008. DOI: 10.1109/MSP.2007.914730
11. N. Ahmed, T. Natarajan and K. R. Rao, “Discrete Cosine Transform,” *IEEE Transactions on Computers* C-23(1):90–93, 1974. DOI: 10.1109/T-C.1974.223784
12. I. E. Sutherland, “Sketchpad: a man-machine graphical communication system,” AFIPS Spring Joint Computer Conference, 1963. DOI: 10.1145/1461551.1461591
13. A. W. Lohmann, R. G. Dorsch, D. Mendlovic, C. Ferreira and Z. Zalevsky, “Space–bandwidth product of optical signals and systems,” *Journal of the Optical Society of America A* 13(3):470, 1996. DOI: 10.1364/JOSAA.13.000470
14. HOLOEYE Photonics, “GAEA-2 10 Megapixel Phase Only Spatial Light Modulator” (4160 × 2464 pixels, 3.74 μm pitch), product page. https://holoeye.com/products/spatial-light-modulators/gaea-2-phase-only/
15. J. W. Goodman, *Introduction to Fourier Optics*, 3rd ed., Roberts & Company, 2005, ch. 5, “Wave-Optics Analysis of Coherent Optical Systems.”
16. C. Jang, K. Bang, G. Li and B. Lee, “Holographic near-eye display with expanded eye-box,” *ACM Transactions on Graphics* 37(6) (SIGGRAPH Asia), 2018. DOI: 10.1145/3272127.3275069
17. S. Choi, B. Chao, J. Yang, M. Gopakumar and G. Wetzstein, “Gaussian Wave Splatting for Computer-Generated Holography,” *ACM Transactions on Graphics* 44(4) (SIGGRAPH), 2025. DOI: 10.1145/3731163; arXiv:2505.06582
18. J. P. Rolland, M. W. Krueger and A. Goon, “Multifocal planes head-mounted displays,” *Applied Optics* 39(19):3209, 2000. DOI: 10.1364/AO.39.003209
19. K. Akeley, S. J. Watt, A. R. Girshick and M. S. Banks, “A stereo display prototype with multiple fo-

cal distances," *ACM Transactions on Graphics* 23(3):804–813, 2004. DOI: 10.1145/1015706.1015804
20. O. Mercier, Y. Sulai, K. Mackenzie, M. Zannoli, J. Hillis, D. Nowrouzezahrai and D. Lanman, "Fast gaze-contingent optimal decompositions for multifocal displays," *ACM Transactions on Graphics* 36(6) (SIGGRAPH Asia), 2017. DOI: 10.1145/3130800.3130846
21. Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart and A. Anandkumar, "Fourier Neural Operator for Parametric Partial Differential Equations," ICLR 2021. arXiv:2010.08895
22. N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio and A. Courville, "On the Spectral Bias of Neural Networks," ICML 2019. arXiv:1806.08734
23. M. Tancik, P. P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron and R. Ng, "Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains," NeurIPS 2020. arXiv:2006.10739
24. L. Gueguen, A. Sergeev, B. Kadlec, R. Liu and J. Yosinski, "Faster Neural Networks Straight from JPEG," NeurIPS 2018. https://papers.nips.cc/paper/2018/hash/7af6266cc52234b5aa339b16695f7fc4-Abstract.html
25. K. Xu, M. Qin, F. Sun, Y. Wang, Y.-K. Chen and F. Ren, "Learning in the Frequency Domain," CVPR 2020, pp. 1737–1746. DOI: 10.1109/CVPR42600.2020.00181
26. TechCrunch, "Holographic startup Envisics partners with Panasonic to fast-track in-car AR tech," 8 January 2021. https://techcrunch.com/?p=2093997

Provenance

Drawn from `wavefield-gate/GATE.md`, `wavefield-gate/exp01_basis_budget.py`, `wavefield-gate/exp01_results.json`, `wavefield-gate/etendue.py`, `wavefield-gate/round3_ops.py`, `docs/SOUNDNESS-AXIS.md` and `docs/SUBSTRATE-GATE.md` in the Quantum-Intelligence research repository.

Published by NeuroGenesys (AW3 Technology, Inc.). Canonical version: https://neurogenesys.si/research/wavefield-gate