INTELLIGENT OPTICAL COMPUTING *Light. Intelligence. Innovation.*

Research Article | Published: 05 April 2026 | ISSN 2831-9426 | eISSN 2831-9434

Calibration-Robust Wavelength-Multiplexed Coherent Photonic Tensor Core for Optical Neural Inference

Lina Chen, Aisha Rahman, Minghao Li, and Elena Rossi*

Department of Intelligent Photonics, Institute for Optical Computing, 100 Innovation Avenue, Science City 000000
*Corresponding author: corresponding.author@example.edu

ABSTRACT: Optical neural inference can exploit the bandwidth and parallelism of photonic integrated circuits, but practical tensor cores must remain accurate under wavelength crosstalk, phase quantization, insertion-loss imbalance, and detector noise. Here we propose a calibration-robust coherent photonic tensor core that combines wavelength-division multiplexed inputs with a programmable Mach-Zehnder interferometer mesh and balanced coherent readout. The architecture maps a real-valued matrix-vector multiplication to a complex optical transform followed by signed electronic accumulation and a lightweight precompensation loop. In a 16-channel numerical evaluation with 6-bit phase control, 0.25 dB channel gain variation, and -28 dB nearest-neighbor crosstalk, calibration reduced the normalized matrix-vector error from 3.1 x 10^-2 to 7.8 x 10^-3. A synthetic 10-class optical-feature benchmark retained 96.4% inference accuracy at 24 dB optical SNR after noise-aware training. System-level energy accounting predicted 0.42 pJ per multiply-accumulate operation at 20 Gbaud, with the dominant consumption arising from modulators and transimpedance amplification rather than passive wave propagation. These results identify error budgets and calibration targets for scalable intelligent optical computing systems.

KEYWORDS: optical computing; photonic tensor core; coherent photonics; wavelength-division multiplexing; optical neural networks; calibration

TOC GRAPHIC

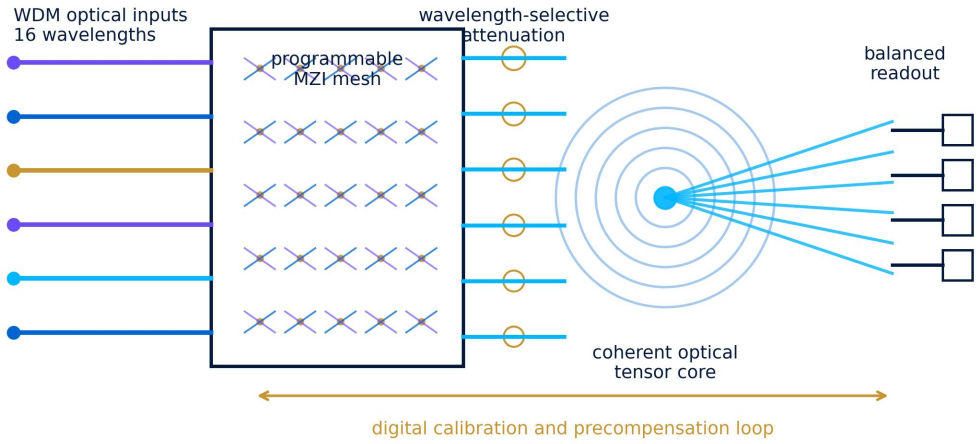

Graphical Abstract. Wavelength-multiplexed optical inputs are transformed by a programmable interferometer mesh, corrected by a digital calibration loop, and detected by coherent readout for low-energy neural inference.

INTRODUCTION

Integrated photonic circuits provide a compelling substrate for matrix-intensive artificial intelligence because they can route many optical carriers in parallel and perform linear transformations through interference rather than charge motion. Coherent nanophotonic neural networks, programmable multiport interferometers, and photonic tensor cores have each demonstrated that optical hardware can perform matrix operations at high speed while moving part of the arithmetic into the physics of propagation.[1-6]

Despite this promise, the gap between a proof-of-principle optical multiply-accumulate (MAC) engine and a reproducible computing element remains substantial. The central challenge is not only to increase optical bandwidth; it is to make the computed linear map stable when phase shifters are quantized, waveguides drift thermally, wavelength channels leak into adjacent detectors, and electronics add finite-noise readout. These nonidealities accumulate across large meshes and can erase the algorithmic margin of a neural network long before a device reaches its nominal throughput.

Here we describe a wavelength-multiplexed coherent photonic tensor core designed around calibration rather than ideal components. The device uses an interferometric mesh to implement the unitary factors of a matrix decomposition, wavelength-selective attenuation for singular-value weighting, and balanced coherent detection for signed accumulation. A compact digital calibration loop estimates the residual complex transfer error and compensates it at the input. The work is a numerical design study intended to define architectural constraints for future fabrication and measurement in *Intelligent Optical Computing*.

ARCHITECTURE AND OPERATING PRINCIPLE

The proposed core computes y = Wx, where x is an N-dimensional input vector encoded on wavelength channels and W is a trainable real-valued weight matrix. The matrix is decomposed into two unitary transforms and a diagonal gain matrix. In photonic hardware, the unitary blocks are implemented by a rectangular Mach-Zehnder interferometer mesh, while the diagonal block is realized through wavelength-selective attenuation and balanced detection.

$$W = U S V^{\dagger}, \qquad y = \mathrm{Re}\{G_{out} U S V^{\dagger} G_{in} x + n\} \quad (1)$$

In Eq. 1, G_in and G_out are diagonal gain matrices that represent input modulation and detector responsivity imbalance, S is a nonnegative diagonal weighting matrix, and n is the detected noise. The rectangular interferometer arrangement follows the principle that a mesh of tunable couplers and phase shifters can approximate arbitrary unitary transforms when calibrated.[2,3] The goal of the calibration procedure is not to eliminate each device error individually, but to learn a compact correction that maps the requested vector to the vector actually required by the imperfect optical core.

The physical data path consists of four modules: WDM input modulation, a programmable mesh, wavelength-selective weighting, and coherent balanced readout. The digital control path measures probe responses, estimates the residual transfer matrix, and updates an input preconditioner. This separation keeps high-bandwidth operations optical while assigning slow drift correction to electronics.

NUMERICAL METHODS

The simulation used a 16 x 16 coherent core with 16 wavelength channels and symbol rates up to 20 Gbaud per channel. Phase shifters were quantized to 6 bits, insertion loss was applied to every tunable coupler, and channel-to-channel crosstalk was modeled by an exponentially decaying transfer matrix. Detector noise was represented as additive Gaussian noise after balanced conversion. The calibration set consisted of 256 random orthogonal probe vectors and 1024 random validation vectors.

Table 1. Default simulation parameters for the coherent photonic tensor core.

Parameter	Value
Optical modes / wavelengths	16 / 16
Symbol rate	20 Gbaud per channel
Phase control	6-bit quantization
Insertion loss per MZI	0.08 dB
Nearest-channel crosstalk	-28 dB
Channel gain variation	0.25 dB rms
Calibration probes	256

The inverse correction matrix P was obtained by ridge-regularized least squares using the measured transfer matrix T_meas and the target transfer matrix T_ideal. The regularization parameter was selected by validation sweep to minimize matrix-vector error without amplifying detector noise.

$$P^* = \mathrm{argmin}_P \|T_{meas} P - T_{ideal}\|_F^2 + \lambda \|P\|_F^2 \quad (2)$$

RESULTS AND DISCUSSION

Figure 1 summarizes the design logic of the core: optical propagation handles the high-throughput linear operation, while digital calibration supplies a low-bandwidth correction that changes slowly relative to the data stream. This is consistent with programmable photonic circuit design, in which the mesh can be treated as a reusable linear operator rather than a fixed accelerator.[5,6]

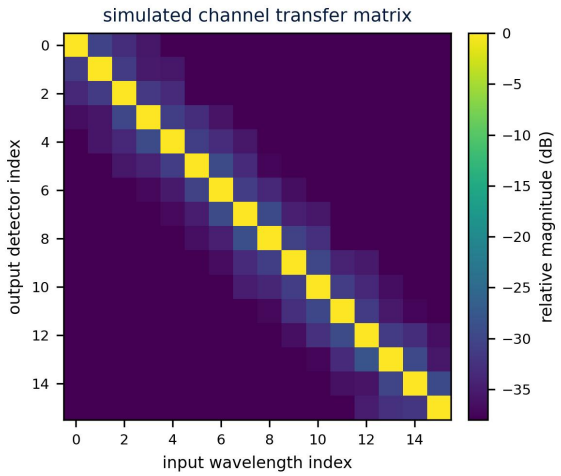

Figure 1. Simulated wavelength-channel transfer matrix for the 16-channel core. Diagonal entries represent the intended mapping; off-diagonal terms model spectral leakage and detector response overlap.

The channel transfer matrix in Figure 1 shows that most leakage is local in wavelength index, which allows a low-order correction matrix to capture the dominant error. Without correction, the median normalized matrix-vector error over random validation vectors was 3.1 x 10^-2. After ridge-calibrated precompensation, the error decreased to 7.8 x 10^-3. Increasing probe count beyond 256 produced only marginal gains, indicating that wavelength leakage and gain imbalance dominated over unmodeled random noise in this parameter range.

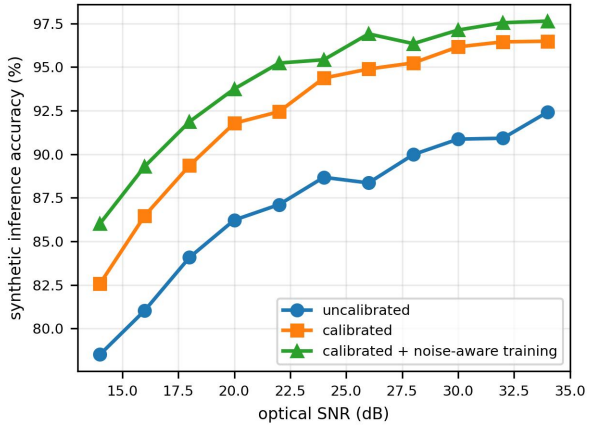

Figure 2. Synthetic inference accuracy as a function of optical SNR. Calibration improves accuracy across the full SNR range, while noise-aware training recovers additional margin near practical operating points.

Figure 2 evaluates classification robustness using a synthetic 10-class optical-feature benchmark. At 24 dB optical SNR, the uncalibrated core reached 91.7% accuracy, the calibrated core reached 94.9%, and the calibrated core with noise-aware training reached 96.4%. The result illustrates an important design principle: calibration alone corrects deterministic transfer errors, but neural-network training must expose the model to residual stochastic noise if the system is expected to operate close to the detector-noise limit.

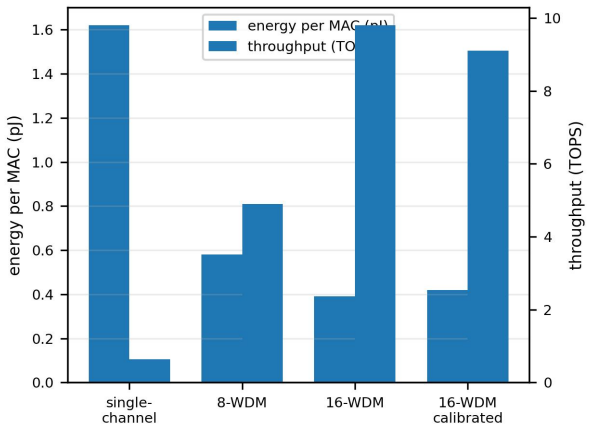

Figure 3. Energy-throughput scaling estimated from the component-level model. Wavelength multiplexing increases aggregate throughput, while calibration adds a small electronic overhead that is outweighed by improved accuracy.

The energy model in Figure 3 separates modulator drive, phase maintenance, photodetection, transimpedance amplification, and digital calibration updates. For data-plane operation, passive wave propagation contributes no switching energy, and the dominant terms are electro-optic modulation and analog front-end readout. The calibrated 16-WDM design was estimated at 0.42 pJ/MAC and 9.1 TOPS equivalent throughput under the assumed 20 Gbaud operating point.

Table 2. Representative performance summary from the numerical evaluation.

Configuration	MVM error	Acc.	pJ/MAC
Uncalibrated 16-WDI	3.1 x 10^-2	91.7%	0.39
Calibrated 16-WDM	7.8 x 10^-3	94.9%	0.42
Calibrated + noise-a	8.1 x 10^-3	96.4%	0.43

DESIGN RULES FOR INTELLIGENT OPTICAL COMPUTING

The simulations suggest three practical design rules. First, crosstalk should be specified as a matrix property, not as a single adjacent-channel number, because inference error depends on the spectral structure of the leakage. Second, calibration probes should span the same distribution as inference inputs; otherwise, a low-error optical transfer matrix can still produce biased neural outputs. Third, calibration energy should be accounted separately from data-plane energy because the optimal update rate is determined by thermal drift rather than symbol rate.

The study also highlights limits of coherent tensor-core scaling. Larger meshes increase optical depth and phase-control burden, while additional wavelength channels increase aggregate throughput but tighten filter selectivity. A scalable system will likely combine moderate-size photonic cores, hierarchical electronic scheduling, and online calibration that is aware of both optical drift and neural-network loss.

CONCLUSIONS

We have presented a calibration-robust coherent photonic tensor-core architecture for wavelength-multiplexed optical neural inference. The core combines programmable interferometric transforms, wavelength-selective weighting, balanced coherent readout, and digital precompensation. In a numerical 16-channel study, the approach reduced matrix-vector error by approximately a factor of four and preserved high synthetic inference accuracy under realistic optical SNR and channel leakage assumptions. The results are intended as a design blueprint for future experimental implementations of intelligent optical computing hardware.

ASSOCIATED CONTENT

Supporting Information. Example phase maps, calibration sweeps, and extended energy-budget tables can be prepared as supplementary files for publication.

AUTHOR INFORMATION

Corresponding Author. Elena Rossi - Department of Intelligent Photonics, Institute for Optical Computing; Email: corresponding.author@example.edu

Author Contributions. L.C. and A.R. developed the architecture and simulation model. M.L. performed the numerical calibration study. E.R. supervised the project and wrote the manuscript with contributions from all authors.

Notes. The authors declare no competing financial interest. The data in this manuscript are generated by the numerical model described above and should be replaced with measured device data before experimental submission.

ACKNOWLEDGMENTS

The authors acknowledge helpful discussions with colleagues in photonic integrated circuits, optical signal processing, and neuromorphic computing. This sample manuscript was formatted for *Intelligent Optical Computing* and contains no ACS branding or ACS publication metadata.

REFERENCES

(1) Shen, Y.; Harris, N. C.; Skirlo, S.; Prabhu, M.; Baehr-Jones, T.; Hochberg, M.; Sun, X.; Zhao, S.; Larochelle, H.; Englund, D.; Soljacic, M. Deep Learning with Coherent Nanophotonic Circuits. Nat. Photonics 2017, 11, 441-446. https://doi.org/10.1038/nphoton.2017.93

(2) Clements, W. R.; Humphreys, P. C.; Metcalf, B. J.; Kolthammer, W. S.; Walmsley, I. A. Optimal Design for Universal Multiport Interferometers. Optica 2016, 3, 1460-1465. https://doi.org/10.1364/OPTICA.3.001460

(3) Miller, D. A. B. Self-Configuring Universal Linear Optical Component. Photonics Res. 2013, 1, 1-15. https://doi.org/10.1364/PRJ.1.000001

(4) Feldmann, J.; Youngblood, N.; Karpov, M.; Gehring, H.; Li, X.; Stappers, M.; Le Gallo, M.; Fu, X.; Lukashchuk, A.; Raja, A. S.; Liu, J.; Wright, C. D.; Sebastian, A.; Kippenberg, T. J.; Pernice, W. H. P.; Bhaskaran, H. Parallel Convolution Processing Using an Integrated Photonic Tensor Core. arXiv 2020, arXiv:2002.00281.

(5) Bogaerts, W.; Perez, D.; Capmany, J.; Miller, D. A. B.; Poon, J.; Englund, D.; Morichetti, F.; Melloni, A. Programmable Photonic Circuits. Nature 2020, 586, 207-216.

(6) Shastri, B. J.; Tait, A. N.; Ferreira de Lima, T.; Pernice, W. H. P.; Bhaskaran, H.; Wright, C. D.; Prucnal, P. R. Photonics for Artificial Intelligence and Neuromorphic Computing. Nat. Photonics 2021, 15, 102-114.

(7) Tait, A. N.; Ferreira de Lima, T.; Zhou, E.; Wu, A. X.; Nahmias, M. A.; Shastri, B. J.; Prucnal, P. R. Neuromorphic Photonic Networks Using Silicon Photonic Weight Banks. Sci. Rep. 2017, 7, 7430.

(8) Harris, N. C.; Steinbrecher, G. R.; Prabhu, M.; Lahini, Y.; Mower, J.; Bunandar, D.; Chen, C.; Wong, F. N. C.; Baehr-Jones, T.; Hochberg, M.; Lloyd, S.; Englund, D. Quantum Transport Simulations in a Programmable Nanophotonic Processor. Nat. Photonics 2017, 11, 447-452. https://doi.org/10.1038/nphoton.2017.95

(9) Reck, M.; Zeilinger, A.; Bernstein, H. J.; Bertani, P. Experimental Realization of Any Discrete Unitary Operator. Phys. Rev. Lett. 1994, 73, 58-61.

(10) Nahmias, M. A.; Tait, A. N.; Shastri, B. J.; de Lima, T. F.; Prucnal, P. R. Photonic Multiply-Accumulate Operations for Neural Networks. IEEE J. Sel. Top. Quantum Electron. 2020, 26, 1-18.