Supplementary Information for

Fully multiplexed photonic tensor computing

Aolong Sun[1,2,6], Junhao Zhao[1,4,6], Fangchen Hu[3,6], Sizhe Xing[1], Yuqin Yuan[1], Jialin He[1], Yongzhu Hu[1], Xuyu Deng[1], Yinjun Liu[1], Ouhan Huang[1], Baiheng Zhao[2], Hancheng Liu[2], Tian Dong[2], Jingkai Zhou[2], Haoyang Sun[2], Liang Chen[5], Chao Shen[1], Feng Bao[1], Ziwei Li[1], Jianyang Shi[3], Wei Chu[3], Bowei Dong[2,*], Nan Chi[1,*], and Junwen Zhang[1,*]

[1] College of Future Information Technology, Fudan University, Shanghai, China

[2] Institute of Microelectronics, Agency for Science, Technology and Research, Singapore

[3] Zhangjiang Laboratory, Shanghai, China

[4] Shanghai Innovation Institute, Shanghai, China

[5] State Key Laboratory of Multimedia Information Processing, School of Computer Science, Peking University, Beijing, China

[6] These authors contributed equally to this work: Aolong Sun, Junhao Zhao, Fangchen Hu

[*] Corresponding authors: junwenzhang@fudan.edu.cn, nanchi@fudan.edu.cn, dong_bowei@ime.a-star.edu.sg

Contents

Supplementary Note 1: Physical framework of fully multiplexed photonic tensor computing

We formulate fully multiplexed photonic tensor computing by linking the discrete tensor representation of the input data to its continuous-time signal description and the underlying multiplexed optical field. Let $\mathcal{X} \in \mathbb{R}^{R\times N_f\times N_\lambda\times N_m\times T}$ denote the input tensor, whose elements are written as $x_{r,n,\ell,m,\tau}$. Here, K and R denote the numbers of output and input ports, respectively; N_f, N_λ, N_m, and T denote the numbers of RF subcarriers, wavelengths, modes, and time symbols, respectively; and n, ℓ, m, and τ index the corresponding tensor slices. Let $\hat{x}_{r,n,\ell,m}(t)$ denote the continuous-time complex envelope associated with the discrete sequence $x_{r,n,\ell,m,\tau}$, sampled at symbol period T_s. For the r-th input port, the ℓ-th wavelength channel, and the m-th mode, the RF-multiplexed electrical envelope is written as

$$s_{r,\ell,m}(t) = \sum_{n=1}^{N_f} \hat{x}_{r,n,\ell,m}(t) e^{j\Omega_n t}, \Omega_n = 2\pi f_n, \tag{S1.1}$$

where f_n is the frequency of the n-th RF subcarrier. Considering linear intensity modulation, the modulated optical envelope $A_{r,l,m}(t)$ is defined such that its instantaneous optical intensity follows the applied waveform: $\left|A_{r,l,m}(t)\right|^2 = \bar{P}_{r,l,m} + \gamma_{r,l,m}\,\Re\{s_{r,l,m}(t)\}$, where $\bar{P}_{r,l,m}$ denotes the average optical power and $\gamma_{r,l,m}$ is the electro-optic conversion coefficient.

In the present architecture, the wavelength dimension plays a dual role. In the main text, the wavelength dimension is represented by the channel index ℓ for conceptual clarity. In the implementation-level derivation here, each wavelength channel ℓ physically corresponds to a wavelength group $\{\omega_{\ell,r}\}$, in which the r-th optical carrier $\omega_{\ell,r}$ is assigned to the r-th input port. In this way, the channel index ℓ provides inter-group WDM scalability, whereas the intra-group carriers $\omega_{\ell,r}$ enable interference-free optical accumulation across the input ports. After wavelength multiplexing, the signals indexed by different mode channels are first carried by the on-chip fundamental mode and are physically routed in different input waveguides. Therefore, for the m-th mode channel, the wavelength-multiplexed field at the r-th input port can be written as

$$E_{r,m}(\boldsymbol{\rho},t) = \sum_{\ell=1}^{N_\lambda} A_{r,\ell,\mathrm{m}}(t)\, \psi(\boldsymbol{\rho};\omega_{\ell,r})\, e^{j\omega_{\ell,r}t}, \tag{S1.2}$$

After subsequent mode multiplexing, these channels are converted to their corresponding guided modes, yielding the fully multiplexed field

$$E_r(\boldsymbol{\rho},t) = \sum_{m=1}^{N_m} \sum_{\ell=1}^{N_\lambda} A_{r,\ell,m}(t)\, \psi_m(\boldsymbol{\rho};\omega_{\ell,r})\, e^{j\omega_{\ell,r}t}, \tag{S1.3}$$

where $\boldsymbol{\rho}$ denotes the transverse spatial coordinate, $\psi_m(\boldsymbol{\rho};\omega_{\ell,r})$ is the transverse profile of the m-th guided mode.

The programmed computation performed by the FieldCore is described by the weight matrix $\boldsymbol{\mathcal{W}} = [w_{k,r}] \in \mathbb{R}^{K\times R}$, where $w_{k,r}$ denotes the kernel weight from the r-th input port to the k-th output port. The optical field at the k-th output is therefore

$$E_k(\boldsymbol{\rho},t) = \sum_{r=1}^{R} \sqrt{w_{k,r}}\, E_r(\boldsymbol{\rho},t). \tag{S1.4}$$

Substituting the input-field expansion yields

$$E_k(\boldsymbol{\rho},t) = \sum_{r=1}^{R} \sum_{m=1}^{N_m} \sum_{\ell=1}^{N_\lambda} \sqrt{w_{k,r}}\, A_{r,\ell,m}(t)\, \psi_m(\boldsymbol{\rho};\omega_{\ell,r})\, e^{j\omega_{\ell,r}t}. \tag{S1.5}$$

After the computation, mode demultiplexing is first applied. Selecting the m-th guided mode gives

$$E_{k,m}(\boldsymbol{\rho},t) = \sum_{r=1}^{R} \sum_{\ell=1}^{N_\lambda} \sqrt{w_{k,r}}\, A_{r,\ell,m}(t)\, \psi_m(\boldsymbol{\rho};\omega_{\ell,r})\, e^{j\omega_{\ell,r}t}. \tag{S1.6}$$

Wavelength demultiplexing is then performed. Selecting the ℓ-th wavelength group yields

$$E_{k,\ell,m}(\boldsymbol{\rho},t) = \sum_{r=1}^{R} \sqrt{w_{k,r}}\, A_{r,\ell,m}(t)\, \psi_m(\boldsymbol{\rho};\omega_{\ell,r})\, e^{j\omega_{\ell,r}t}. \tag{S1.7}$$

The field associated with the selected optical channel can therefore be represented by a scalar complex amplitude

$$\mathcal{E}_{k,\ell,m}(t) = \sum_{r=1}^{R} \sqrt{w_{k,r}}\, A_{r,\ell,m}(t)\, e^{j\omega_{\ell,r}t}. \tag{S1.8}$$

The detected photocurrent is governed by square-law detection, $i_{k,\ell,m}(t) \propto \left|\mathcal{E}_{k,\ell,m}(t)\right|^2$. Expanding the square gives

$$i_{k,\ell,m}(t) \propto \sum_{r=1}^{R} w_{k,r} \left|A_{r,\ell,m}(t)\right|^2 + \sum_{r\neq r'} \sqrt{w_{k,r} w_{k,r'}}\, A_{r,\ell,m}(t) A^*_{r',\ell,m}(t) e^{j(\omega_{\ell,r}-\omega_{\ell,r'})t}. \tag{S1.9}$$

The first term is the desired result, whereas the second term corresponds to beating between different intra-group wavelengths. When the intra-group wavelength spacing satisfies $\frac{|\omega_{\ell,r}-\omega_{\ell,r'}|}{2\pi} \gg B_{\mathrm{Rx}}, r \neq r'$, where B_{Rx} is the receiver bandwidth, these inter-wavelength beating terms fall outside the electrical passband and are suppressed. After removing the DC component, the detected in-band electrical signal reduces to $r_{k,\ell,m}(t) \propto \sum_{r=1}^{R} w_{k,r}\, \Re\{s_{r,l,m}(t)\}$. Substituting the FDM composition gives

$$r_{k,\ell,m}(t) \propto \frac{1}{2}\sum_{n=1}^{N_f}\sum_{r=1}^{R} w_{k,r}\left(\hat{x}_{r,n,l,m}(t)e^{j\Omega_n t} + \hat{x}^*_{r,n,l,m}(t)e^{-j\Omega_n t}\right). \quad \text{(S1.10)}$$

Finally, RF subcarrier demultiplexing is performed by down-conversion and low-pass filtering:

$$y_{k,n,\ell,m}(t) = \mathcal{N}\left\{\mathrm{LPF}\left\{r_{k,\ell,m}(t)e^{-j\Omega_n t}\right\}\right\} = \sum_{r=1}^{R} w_{k,r}\,\hat{x}_{r,n,\ell,m}(t). \quad \text{(S1.11)}$$

where $\mathcal{N}\{\cdot\}$ denotes normalization by the proportionality constants, and $\mathrm{LPF}\{\cdot\}$ denotes the low-pass filtering. Sampling at the symbol times gives the recovered discrete output tensor slices,

$$y_{k,n,\ell,m,t} = y_{k,n,\ell,m}(t)\big|_{t=\tau T_s} = \sum_{r=1}^{R} w_{k,r}\,x_{r,n,\ell,m,\tau}. \quad \text{(S1.12)}$$

Therefore, the overall tensor mapping implemented by the FieldCore can be written as

$$\mathcal{Y} \in \mathbb{R}^{K\times N_f\times N_\lambda\times N_m\times T}, y_{k,n,\ell,m,\tau} = \sum_{r=1}^{R} w_{k,r}\,x_{r,n,\ell,m,\tau}. \quad \text{(S1.13)}$$

Supplementary Note 2: Device Design and Characterization

The proposed FieldCore is implemented as a photonic integrated circuit (PIC) comprising multiple inverse-designed photonic building blocks, including mode multiplexers, optical multipliers, tunable couplers, waveguide crossings, and low-crosstalk bends, all supporting broadband dual-mode (TE_0 and TE_1) operation across the C-band.

A. Structure and simulated performance of the mode multiplexer

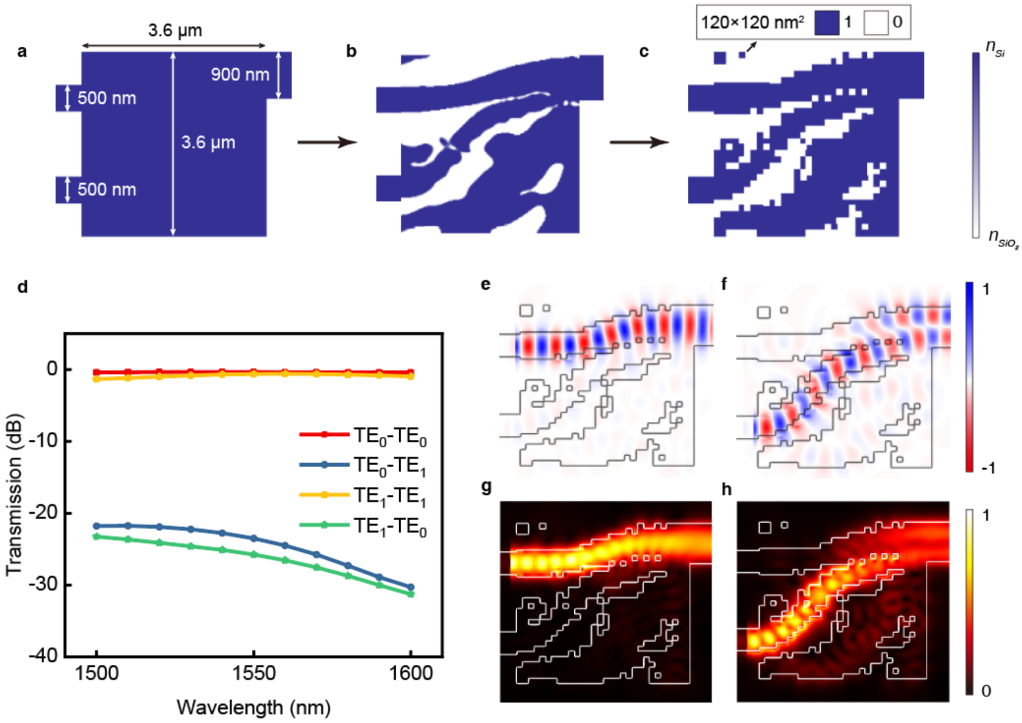

Fig. S1. Design and simulation of the dual-mode multiplexer. (a) Initial refractive index distribution; **(b)** analog optimization result; **(c)** final digitized refractive index pattern; **(d)** Simulated transmission spectra for TE_0 and TE_1 modes; **(e–f)** Simulated E_y field distribution for TE_0 and TE_1 modes at 1550 nm; **(g–h)** Simulated profile of electric field intensity for TE_0 and TE_1 modes at 1550 nm.

The mode multiplexer is optimized using an edge-guided analog-and-digital optimization (EG-ADO) method[1]. The initial structure of the mode multiplexer is shown in Fig. S1(a). The design region occupies a compact footprint of only 3.6 × 3.6 μm^2, with a single-mode input waveguide of 500 nm width and a multimode output waveguide of 900 nm width. During the optimization process based on the EG-ADO algorithm, we first perform adjoint method-based topology optimization to update the refractive index distribution, yielding the intermediate analog structure shown in Fig. S1(b). To suppress the emergence of small features and sharp angles, a circular Heaviside filter with a radius of 240 nm is applied after each iteration. However, the resulting analog pattern remains challenging for fabrication. To address this, we employ an edge-guided analog-to-digital conversion approach: pixels located along the material boundary between silicon and silicon dioxide are subjected to an additional discrete optimization round using direct binary search (DBS) to determine their material assignment, while the remaining pixels are binarized through direct index mapping. The final digital structure, shown in Fig. S1(c), consists of regular square holes with a minimum feature size of 120 × 120 nm^2, providing high fabrication robustness.

The simulated transmission spectra are presented in Fig. S1(d), illustrating the response for TE_0-mode excitation from both the upper and lower ports. For both channels, the insertion loss at 1550 nm is below 0.6 dB, and the crosstalk remains under –22 dB across the full 100 nm wavelength range. Figs. S1(e)-(f) show the distribution of the E_y field component for the two modes, respectively. Figs. S1(g)-(h) display the corresponding electric field intensity distributions, validating the low-loss and low-crosstalk mode conversion process of the device.

B. Structure and simulated performance of the 3-dB splitter

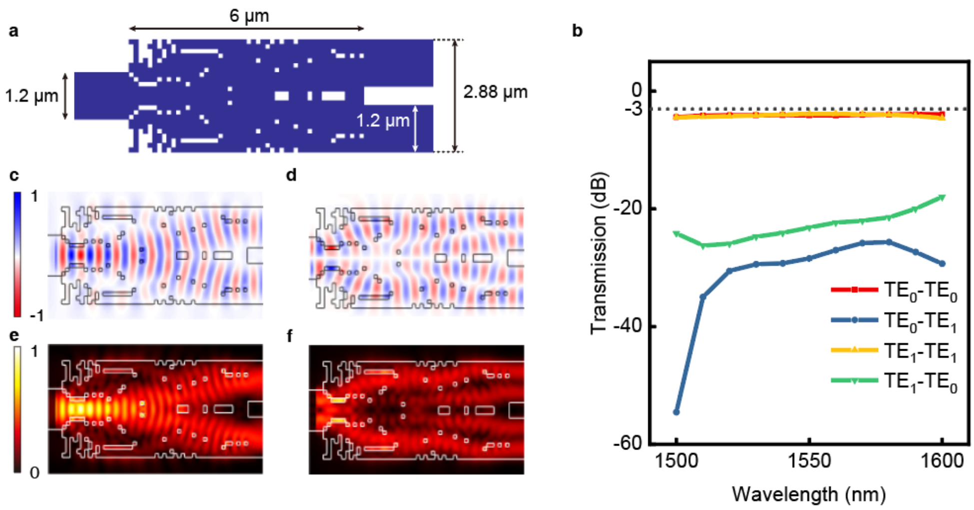

Fig. S2. Design and simulation of the dual-mode 3-dB splitter. (a) Final digitized refractive index distribution; **(b)** Simulated transmission spectra for TE_0 and TE_1 modes; **(c − d)** Simulated E_y field distribution for TE_0 and TE_1 inputs at 1550 nm; **(e − f)** Simulated profile of electric field intensity for TE_0 and TE_1 inputs at 1550 nm.

The dual-mode 3-dB splitter, employed in the optical multiplier, is designed via a hybrid analog-digital method[2]. Fig. S2(a) shows the structure of the inverse-designed dual-mode 3-dB splitter, capable of simultaneously splitting both TE_0 and TE_1 modes with balanced output and minimal crosstalk. The design region spans 6 μm × 2.88 μm and is discretized into square pixels with a minimum feature size of 120 nm. Due to the vertical mirror symmetry of the structure, the simulated optical responses at the upper and lower output ports are identical. As shown in Fig. S2(b), the transmission spectra of both modes remain flat across the 1.5–1.6 μm wavelength band. The insertion losses are less than 1.12 dB, and intermodal crosstalks are suppressed below –23 dB over the entire C-band for the two modes. Figs. S2(c)–(d) present the simulated E_y field component distributions for TE_0 and TE_1 inputs, respectively, while Figs. S2(e)–(f) show the corresponding electric field intensity profiles, illustrating mode-preserving propagation through the splitter.

C. Structure and simulated performance of the crossing

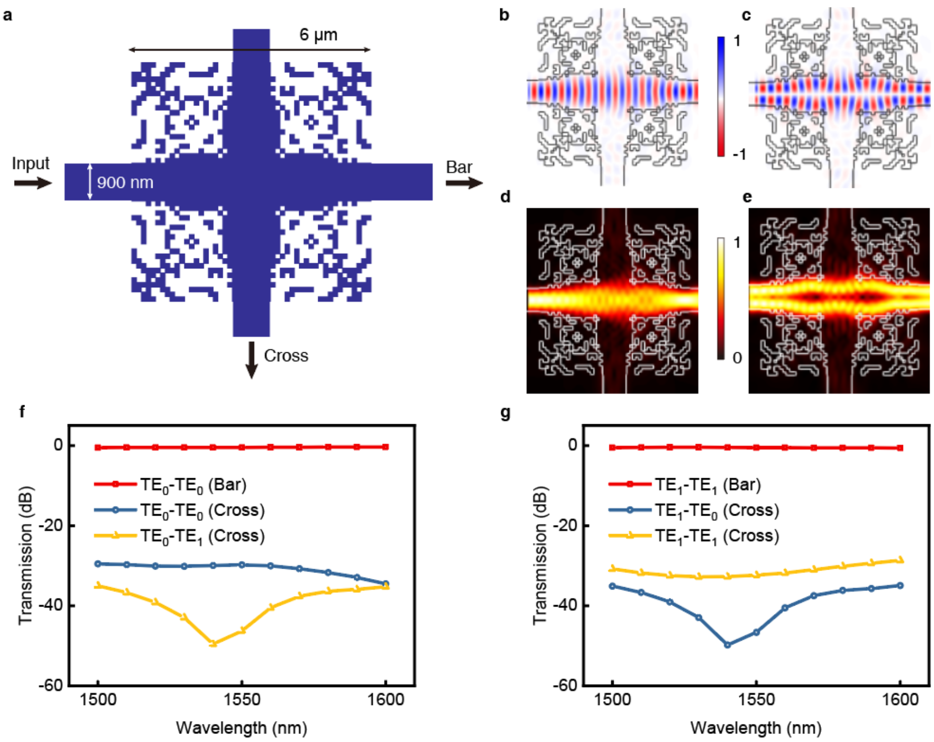

Fig. S3. Design and simulation of the dual-mode crossing. (a) Optimized structure of the crossing; **(b–c)** Simulated E_y field distribution for TE_0 and TE_1 inputs at 1550 nm; **(d–e)** Simulated profiles of electric field intensity for TE_0 and TE_1 inputs at 1550 nm. **(f)** Simulated transmission spectra for TE_0 input; **(g)** Simulated transmission spectra for TE_1 input.

The waveguide crossing is implemented using a DBS algorithm[3], where symmetry constraints are exploited to reduce the computational burden. The final structure of the dual-mode crossing is shown in Fig. S3(a). The device occupies a compact footprint of 6 μm × 6 μm and is discretized into 120 nm × 120 nm square pixels. Figs. S3(b)-(c) depict the distributions of the E_y field component at 1550 nm for TE_0 and TE_1 excitation, respectively, while Fig. S3(d)-(e) present the corresponding electric field intensity profiles. The simulated transmission spectra for TE_0 and TE_1 input modes are shown in Figs. S3(f)-(g), respectively. Due to the mirror symmetry of the device geometry, mode conversion between the symmetric mode (TE_0) and the antisymmetric mode (TE_1) at the bar port is inherently suppressed. Consequently, the crosstalk resulting from this intermodal conversion is negligible, with simulated values below –100 dB. Here, crosstalk is quantified by the power scattered into the cross port. Within the wavelength range of 1500–1600 nm, the insertion loss remains below 0.63 dB, and the crosstalk is suppressed to below –28.3 dB for both modes, demonstrating high modal isolation and low-loss performance.

D. Structure and simulated performance of the multimode waveguide bend

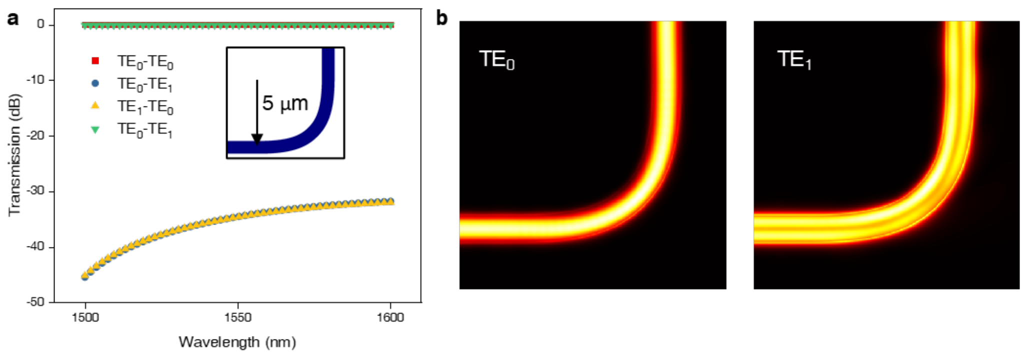

Fig. S4. Design and simulation of the dual-mode bend. (a) Simulated transmission spectra for TE_0 and TE_1 inputs. The inset depicts the bend structure with an effective radius of 5 μm. **(b)** Simulated profiles of electric field intensity for TE_0 and TE_1 inputs at 1550 nm.

To enable compact and low-crosstalk interconnects in our FieldCore, we design a dual-mode waveguide bend based on a quasi-uniform B-spline trajectory, optimized using particle swarm optimization (PSO). As shown in the inset of Fig. S4(a), the bend has an effective radius of 5 μm and supports broadband operation for both TE_0 and TE_1 modes within a 900 nm-wide waveguide. The design method follows a similar procedure to that based on Bezier curves[4], but we replace the conventional Bezier curve with a quasi-uniform B-spline to achieve smoother waveguide geometries and higher-order continuity, which enhances mode preservation during propagation through the bend. The optimization simultaneously targets minimal insertion loss and intermodal crosstalk over the entire C-band.

The simulated transmission spectra in Fig. S4(a) show that the insertion losses remain below 0.009 dB for TE_0 and 0.081 dB for TE_1 modes, while the crosstalk stays below –31 dB throughout the 1500-1600 nm wavelength range. Fig. S4(b) presents the corresponding electric field intensity profiles at 1550 nm, indicating efficient and mode-preserving transmission for both input modes.

E. Structure and simulated performance of the mode-insensitive phase shifter

In the FieldCore, both weight loading and tunable coupling are implemented using Mach-Zehnder interferometer (MZI) architectures. To achieve mode-insensitive operation, we design a mode-insensitive phase shifter (MIPS), as shown in Fig. S5(a). The device consists of a wide multimode silicon waveguide under a titanium nitride (TiN) heater. The waveguide width is set to 4 μm, and the heater width is set to 6 μm to ensure uniform thermal distribution. It is noteworthy that the variation in the effective index with respect to temperature ($\mathrm{d}n_{eff}/\mathrm{d}T$) between different TE modes decreases as the waveguide becomes wider (see Fig. S5(h)). When the waveguide width reaches 4 μm, the difference in $\mathrm{d}n_{eff}/\mathrm{d}T$ between TE_0 and TE_1 modes drops below 1%, enabling uniform phase shifting for both modes. To support dual-mode operation, the input waveguide width is set to 900 nm (see Fig. S5(g)) and is adiabatically tapered to 4 μm. Fig. S5(b) shows the simulated transmission spectra of an MIPS. Figs. S5(e) and S5(f) present the simulated electric field intensity profiles for TE_0 and TE_1 inputs at the wavelength of 1550 nm, respectively, confirming strong mode confinement and minimal intermodal crosstalk.

Furthermore, we simulate the temperature distribution under applied voltage using the ANSYS Lumerical HEAT solver[5]. To reduce heat diffusion into the taper regions where the difference in dn_{eff}/dT between modes becomes more pronounced, the heater is designed to be 4 μm shorter than the waveguide, as illustrated in Fig. S5(c). The cross-sectional temperature profile of the waveguide is shown in Fig. S5(d). Finally, we perform co-simulation using the HEAT and FDTD solvers to extract the phase shifts under different voltages. As shown in Fig. S5(i), the phase shift curves for TE_0 and TE_1 are nearly identical, validating the mode-insensitive behavior.

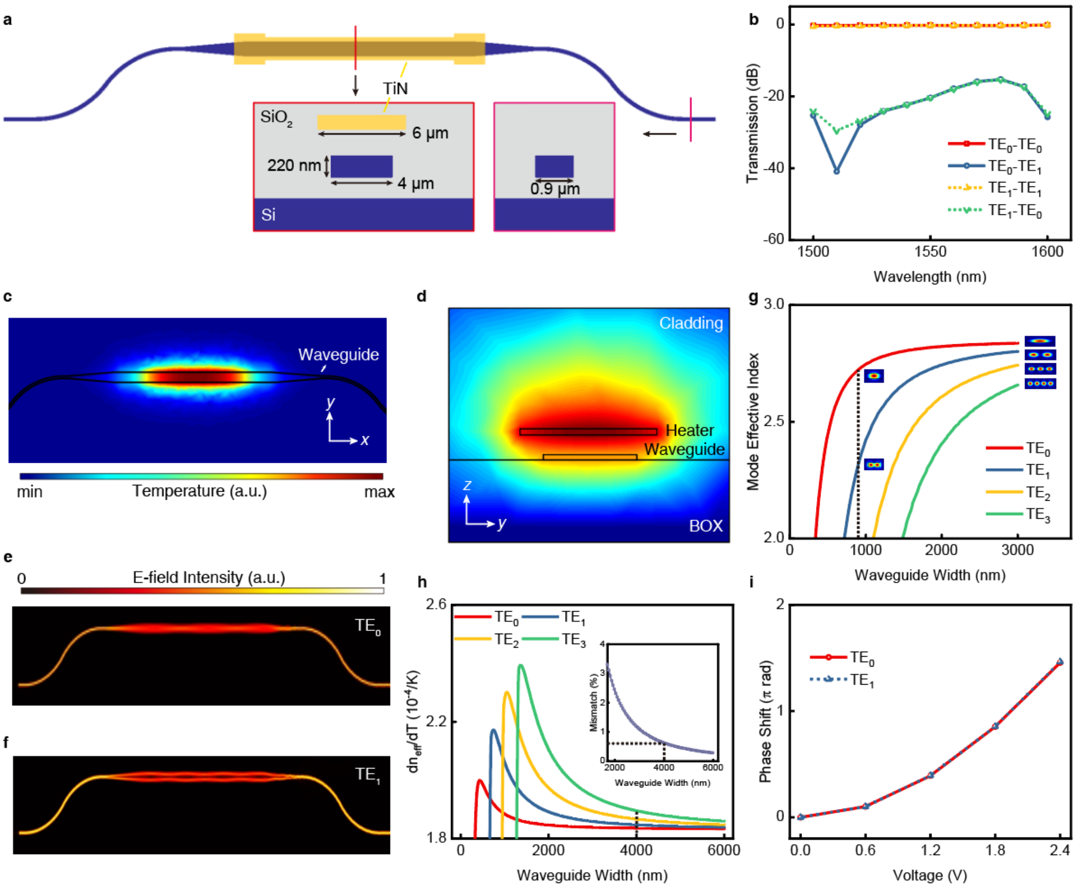

Fig. S5. Design and simulation of the mode-insensitive phase shifter (MIPS). (a) Schematic of the MIPS structure. Insets show cross-sectional views of the heater and waveguide. **(b)** Simulated transmission spectra for TE_0 and TE_1 modes. **(c–d)** Simulated temperature distribution in the MIPS under electrical bias: **(c)** in the *x*-*y* plane, **(d)** in the *y*-*z* plane. **(e–f)** Simulated electric field intensity profiles for TE_0 and TE_1 inputs at 1550 nm. **(g)** Mode effective index for TE_0–TE_3 as a function of waveguide width, showing coexistence of TE_0 and TE_1 at 900 nm width. **(h)** Simulated dn_{eff}/dT for TE_0–TE_3 versus waveguide width; inset shows mismatch between TE_0 and TE_1, which drops below 1% at 4000 nm width. **(i)** Phase shift versus applied voltage for TE_0 and TE_1 modes.

F. Fabrication tolerance analysis for digital metamaterial-based components

The FieldCore integrates three types of inverse-designed digital metamaterial devices: mode MUXs, 3-dB splitters, and dual-mode crossings. These devices are fabricated based on subwavelength pixelated patterns with a foundry-compatible minimum feature size of 120 nm. To evaluate their fabrication robustness, we analyze the impact of pixel size variations on device performance. Specifically, we simulate the insertion loss and intermodal crosstalk under pixel size perturbations ranging from –20 nm to +20 nm.

The simulation results indicate that both the mode multiplexer (see Fig. S6) and the crossing (see Fig. S8) exhibit minimal performance degradation within this tolerance range, while the 3-dB splitter (see Fig. S7) maintains acceptable performance for pixel size variations from –20 nm to +10 nm. According to batch testing data from multiple wafers provided by a mainstream commercial foundry[6], the average variation in waveguide width is approximately ±6.4 nm. Therefore, all three devices demonstrate fabrication robustness to meet practical manufacturing requirements, confirming their potential compatibility with standard foundry processes.

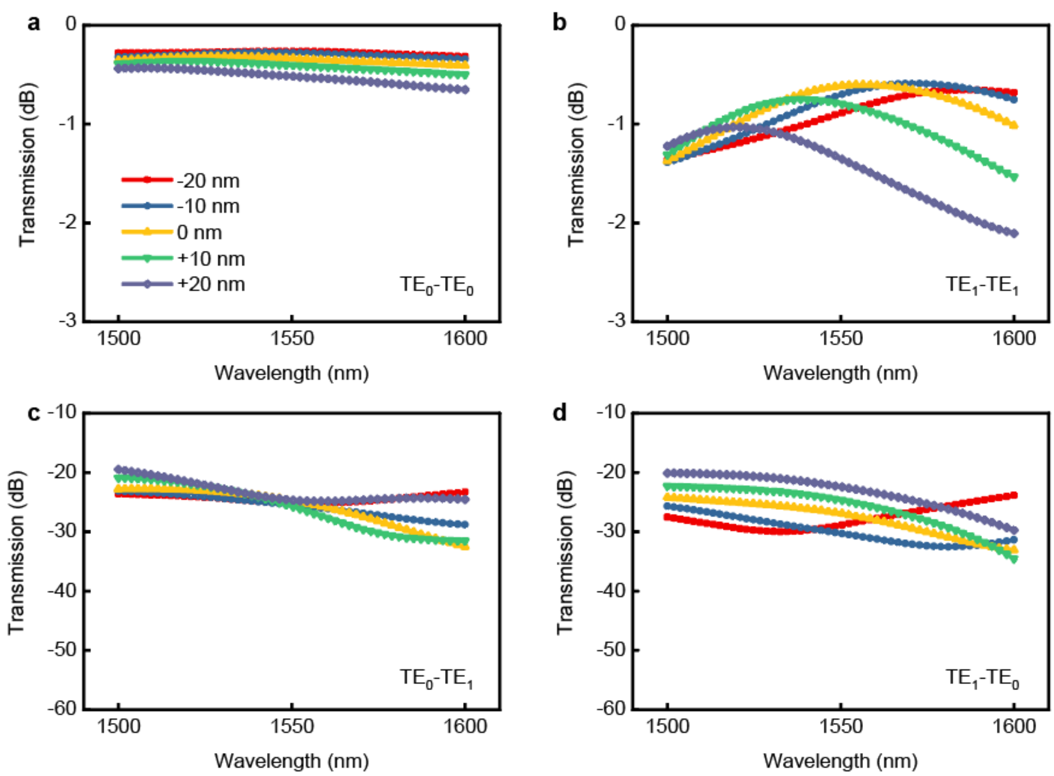

Fig. S6. The simulated transmission spectra of the mode MUX under −20 to + 20 nm pixel size variations. (a-b) Simulated insertion losses for **(a)** TE_0 and **(b)** TE_1 modes. **(c-d)** Simulated intermodal crosstalks for **(c)** TE_0 and **(d)** TE_1 modes.

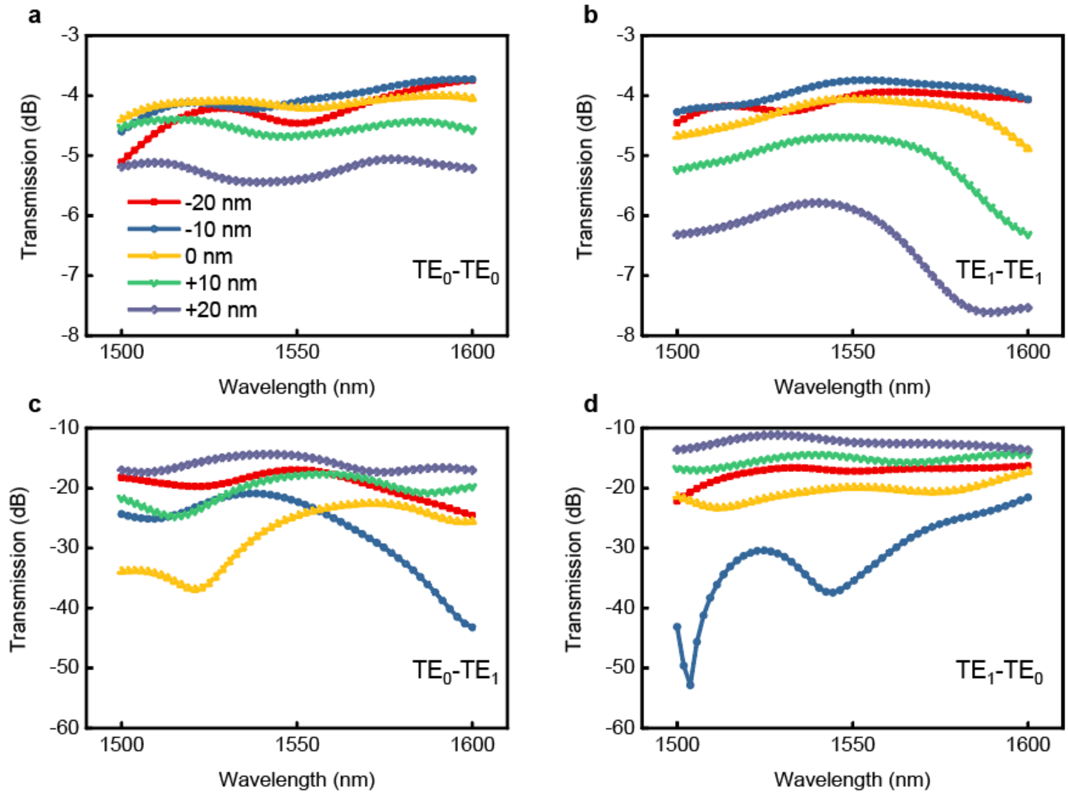

Fig. S7. The simulated transmission spectra of the 3-dB splitter under −20 to + 20 nm pixel size variations. (a-b) Simulated insertion losses for **(a)** TE_0 and **(b)** TE_1 modes. **(c-d)** Simulated intermodal crosstalks for **(c)** TE_0 and **(d)** TE_1 modes.

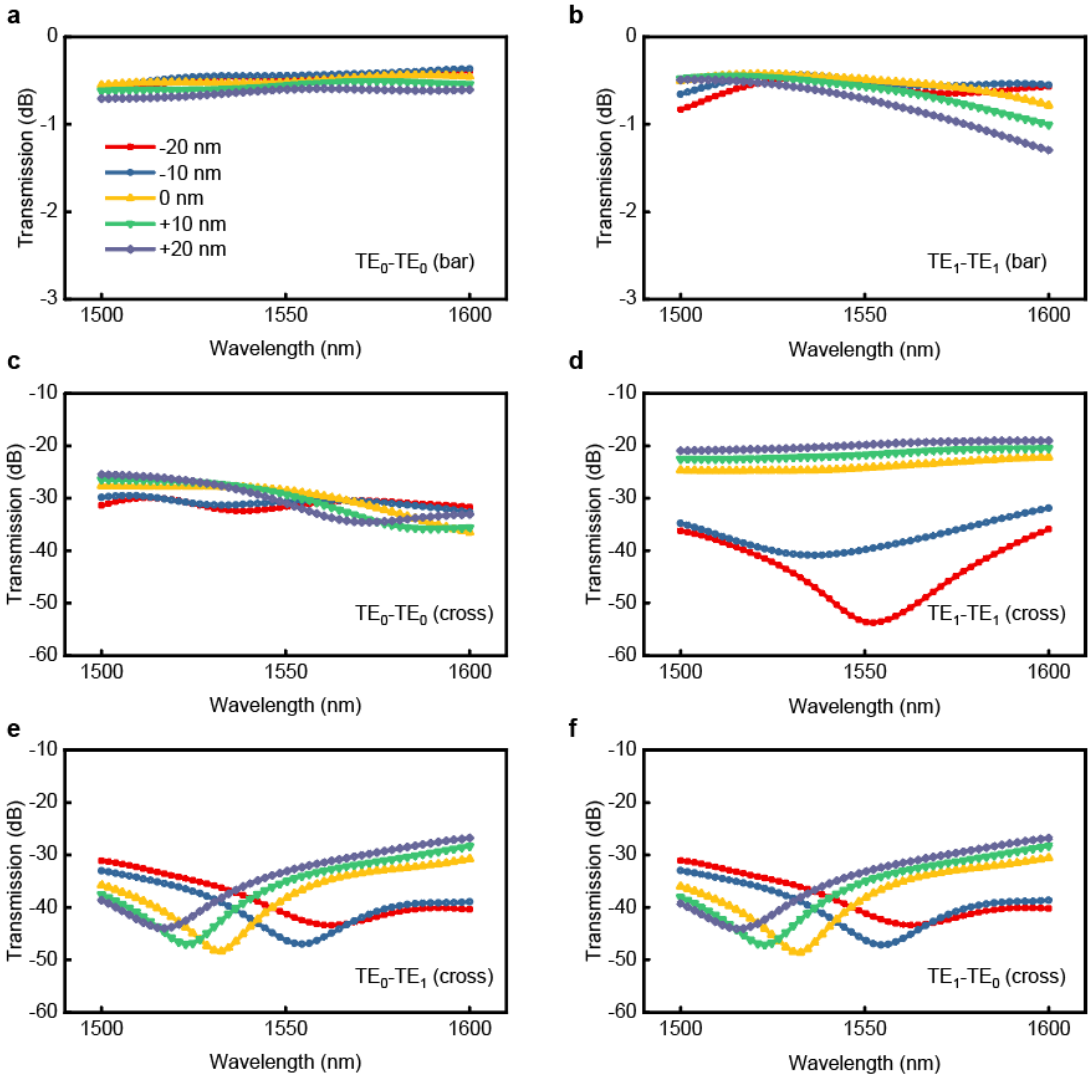

Fig. S8. The simulated transmission spectra of the dual-mode crossing under −20 to + 20 nm pixel size variations. (a-b) Transmission to the bar port for TE_0 and TE_1 inputs, respectively. **(c–f)** Transmission to the cross port for TE_0 and TE_1 inputs, respectively.

G. Experimental measurement results of fabricated devices

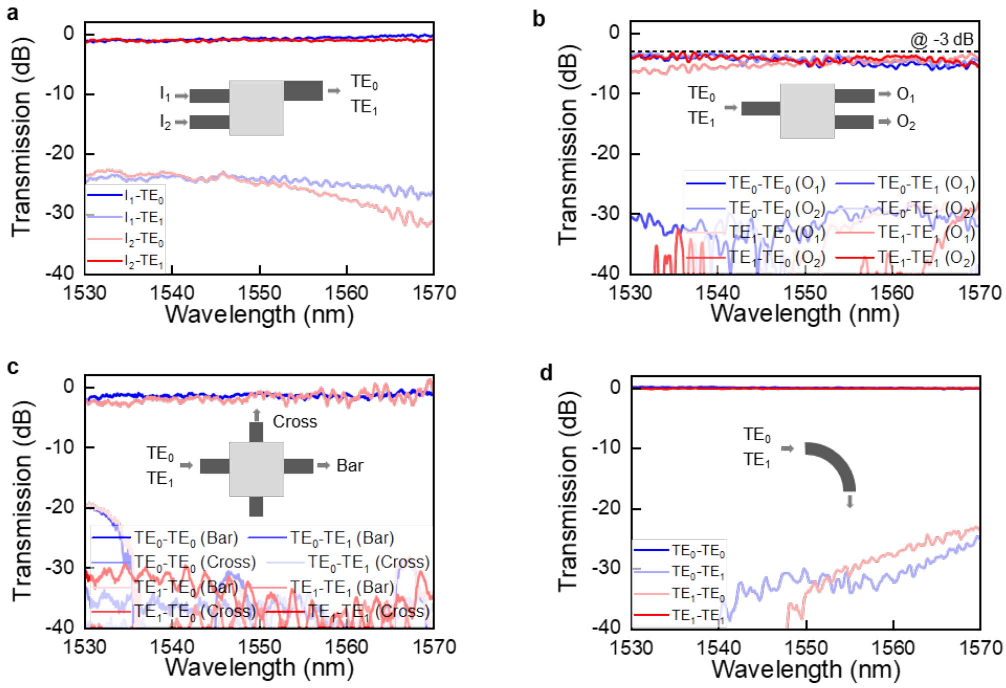

Fig. S9. Measurement results for fabricated inverse-designed components, including (a) mode MUX, **(b)** 3-dB splitter, **(c)** crossing, and **(d)** multimode bend.

In the main text, we present the characterization results of a representative device, namely the optical multiplier. Here, we further provide the measured performance of the inverse-designed unit devices, including the mode MUX, 3-dB splitter, waveguide crossing, and multimode bend (Fig. S9). For the MUX measurement, we fabricate a back-to-back link consisting of a pair of MUX and deMUX. The measured spectra are first normalized by removing the grating-coupler loss, and the response of a single MUX is then extracted by dividing the measured insertion loss by two. For the splitter and crossing, a pair of mode MUX/deMUX is similarly used to perform mode conversion, and the measured spectra are normalized by subtracting the spectrum of the back-to-back MUX link. For the multimode bend measurement, a cascade of ten 90° bends is introduced between the mode MUX and deMUX, and the response of a single bend is then obtained through normalization. The results show that all these inverse-designed unit devices exhibit flat transmission spectra across the entire 40-nm measurement window, supporting operation over the full C band, while maintaining low insertion loss and high mode isolation exceeding 20 dB.

Supplementary Note 3: Verification of Mode- and Wavelength-Domain Operation

A. Signal fidelity under parallel dual-mode transmission

To assess the impact of simultaneous dual-mode operation on signal fidelity, we compare single-mode and dual-mode transmission conditions. An 8-level, 10-GBaud staircase signal is generated by the AWG and applied to the input modulators. In the single-mode case, only TE_0 channel is activated, whereas in the dual-mode case, both TE_0 and TE_1 channels are operated simultaneously.

Fig. S10(a) shows the measured time-domain waveforms received from the TE_0 channel under the single-mode and dual-mode conditions. Over a 200-ns observation window, additional fluctuations are visible when both modes are simultaneously activated. A closer view of the 40-50 ns interval, however, shows that the waveform profile remains largely preserved, with only minor deviations between the two cases. A similar comparison for the TE_1 channel is presented in Fig. S10(b). For the TE_1 channel, slightly larger waveform fluctuations are observed under dual-mode operation, while the overall multi-level signal shape remains preserved. To further quantify this penalty, we extract the signal-to-noise ratio (SNR) from the received waveforms. As shown in Figs. S10(c)-(d), both modes retain relatively high SNRs when the two modes are operated simultaneously. For TE_0, the SNR decreases from 32.13 dB in the single-mode case to 30.16 dB in the dual-mode case, while for TE_1 it decreases from 31.51 dB to 28.8 dB. Although a modest SNR reduction is observed under dual-mode operation, both modes retain relatively high signal fidelity, indicating that simultaneous two-mode transmission remains feasible under the present experimental conditions.

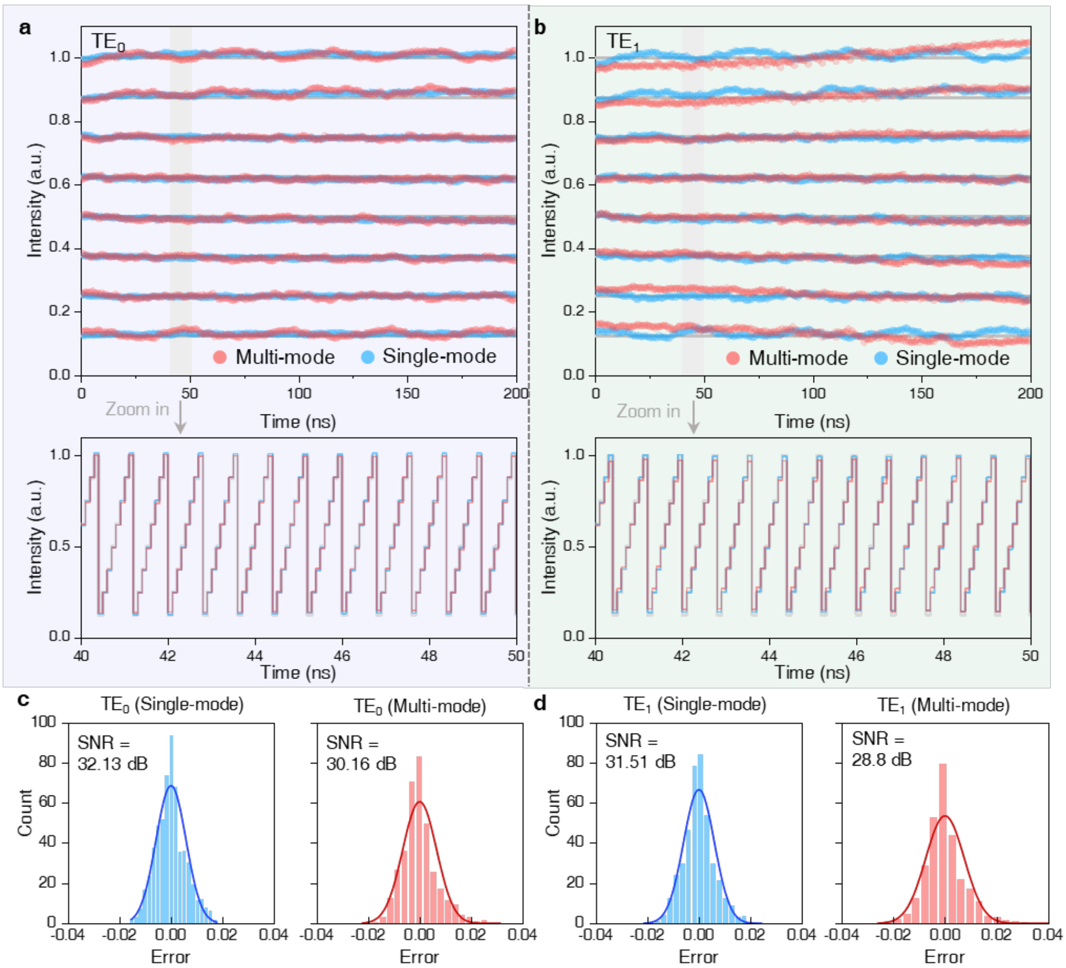

Fig. S10. Signal transmission results for single-mode and multimode operations. (a) Measured temporal waveforms of the TE_0 channel. The top panel shows a 200 ns overview; the bottom panel zooms in on the 40–50 ns interval to highlight waveform fidelity. **(b)** Corresponding temporal waveforms for the TE_1 channel under both conditions. **(c–d)** Histograms of the normalized error distribution for **(c)** TE_0 and **(d)** TE_1, respectively, under single-mode and multimode operation.

B. Wavelength-spacing requirement for WDM-based optical addition

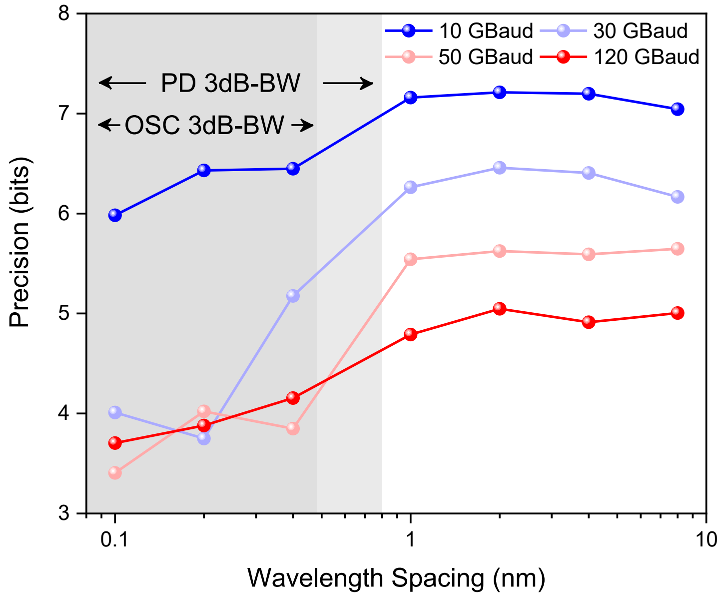

Fig. S11. Influence of wavelength spacing on optical addition accuracy at different baudrates.

In our FieldCore, accurate addition of optical signals relies on the linear addition of optical powers detected by PDs[7,8]. To analyze the wavelength spacing requirement in WDM-based optical addition, we consider the square-law detection behavior of PDs. A more general derivation is provided in Supplementary Note 1. Suppose two optical carriers have electric fields given by: $E_1(t) = A_1(t) \cdot exp(j2\pi f_1 t)$ and $E_2(t) = A_2(t) \cdot exp(j2\pi f_2 t)$. The detected photocurrent can then be expressed as:

$I(t) \propto |E_1(t) + E_2(t)|^2 = A_1(t)^2 + A_2(t)^2 + 2A_1(t)A_2(t) \cdot cos(2\pi(f_1 - f_2)t)$.

The addition of the first two terms is the expected result, while the last term represents a beating component at frequency $|f_1 - f_2|$, which falls within the effective electrical detection bandwidth if the wavelength spacing is too small. This causes distortion in the output signal and reduces addition accuracy. Therefore, accurate optical addition requires a sufficiently large wavelength spacing such that the beat-induced term is strongly attenuated by the receiver response.

We experimentally validate this requirement using the optical addition setup. Since the wavelength-spacing requirement is primarily determined by the electrical detection bandwidth, the TE_0 channel is taken as a representative case. One optical channel is fixed at 1550 nm, while the second channel is swept from 1550.1 nm to 1558 nm. Four signal rates, namely 10 GBaud, 30 GBaud, 50 GBaud, and 120 GBaud, are tested to assess the impact of signal speed on addition precision. As shown in Fig. S11, the calculation precision approaches a plateau once the wavelength spacing exceeds approximately 1 nm for all tested baudrates. This indicates that the beat-induced distortion has been sufficiently suppressed beyond this spacing in the present receiver chain. The required spacing is slightly larger than the nominal bandwidth-based estimate owing to the gradual, rather than abrupt, roll-off of the electrical response.

Supplementary Note 4: Principle and Verification of Frequency-Division Multiplexing

A. Principle of FDM signal generation

A conceptual diagram of the FDM signal generation process is presented in Fig. S12. Here, FDM does not refer to directly encoding information onto the amplitudes of isolated RF tones[7]. Instead, each RF component serves as an electrical subcarrier that carries an up-converted high-speed data stream[9], thereby enabling efficient partitioning and full utilization of the available modulation bandwidth.

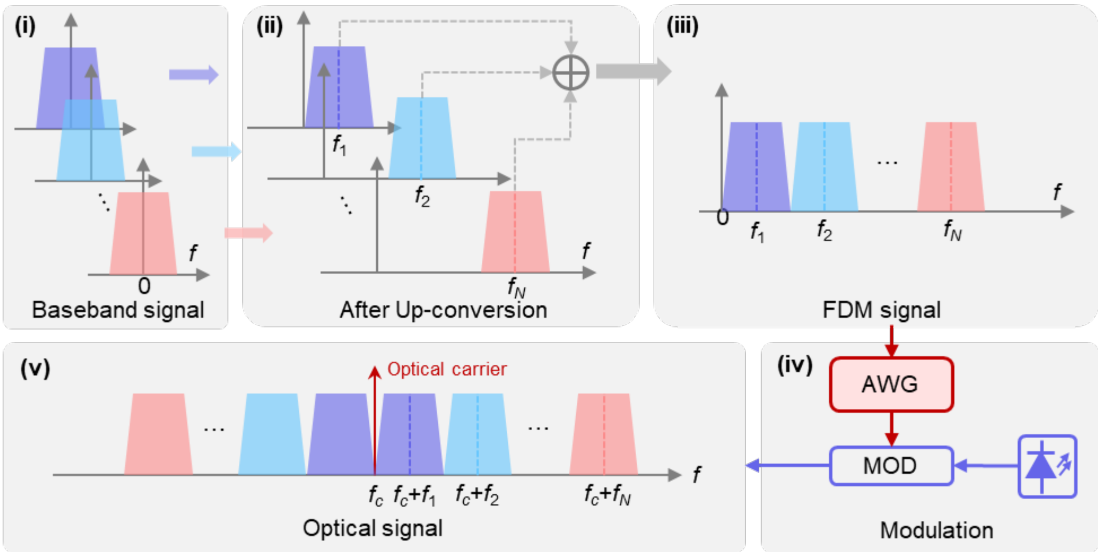

Fig. S12. Principle and implementation flow of frequency-division multiplexing. (i) Independent baseband electrical signals are first generated. **(ii)** Each baseband stream is up-converted to a distinct RF subcarrier at f_1, f_2, …, f_N. **(iii)** The up-converted channels are combined into a composite FDM electrical signal. **(iv)** The waveform is generated by the AWG and used to drive the external modulator. **(v)** The multiplexed RF subcarriers are then mapped onto the optical carrier at f_c, forming the FDM-based optical signal.

B. Process of FDM-based MAC operation

The principle of MAC operations for FDM signal is illustrated in Fig. S13.

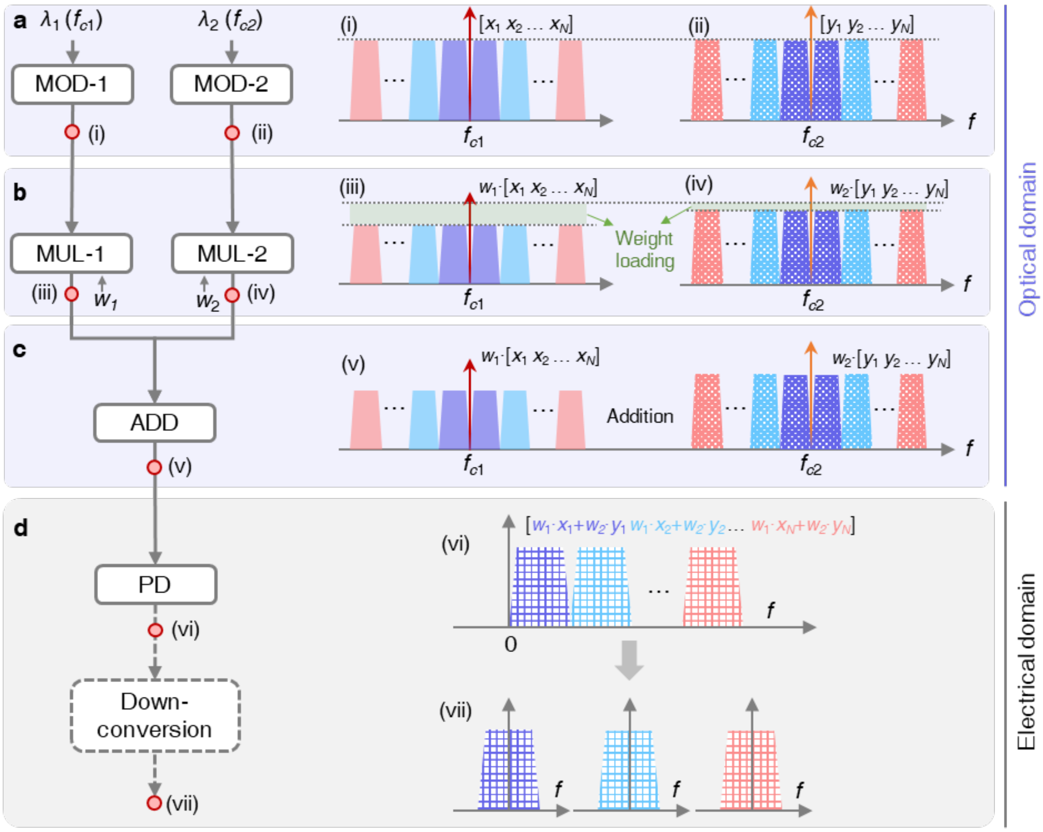

Fig. S13. Principle of FDM-based MAC operation. (a) Two FDM signals are modulated onto two optical carriers at wavelengths λ_1 and λ_2, centered at f_{c1} and f_{c2}, respectively. The two channels share an identical RF subcarrier grid, so that corresponding subcarriers are aligned in radio frequency. **(b)** The two optical signals are routed to separate on-chip multiplier branches (MUL-1 and MUL-2), where weighting coefficients w_1 and w_2 are programmed electrically. Owing to the nearly flat response of the photonic multipliers, all subcarriers carried by the same optical carrier are scaled uniformly by the corresponding weight. **(c)** The weighted optical signals are then combined by the on-chip optical adder. The carrier spacing is chosen to be sufficiently large to avoid sideband overlap and to place the inter-wavelength beating beyond the receiver bandwidth. **(d)** After photodetection, the RF spectra carried by the two optical carriers are mapped onto the same electrical subcarrier grid, while the beat terms between the two wavelengths are rejected by the finite receiver bandwidth. As a result, the n-th subcarrier components from different optical channels add linearly in the electrical domain. Each subcarrier is then down-converted and demodulated individually, thereby completing the parallel FDM MAC operation.

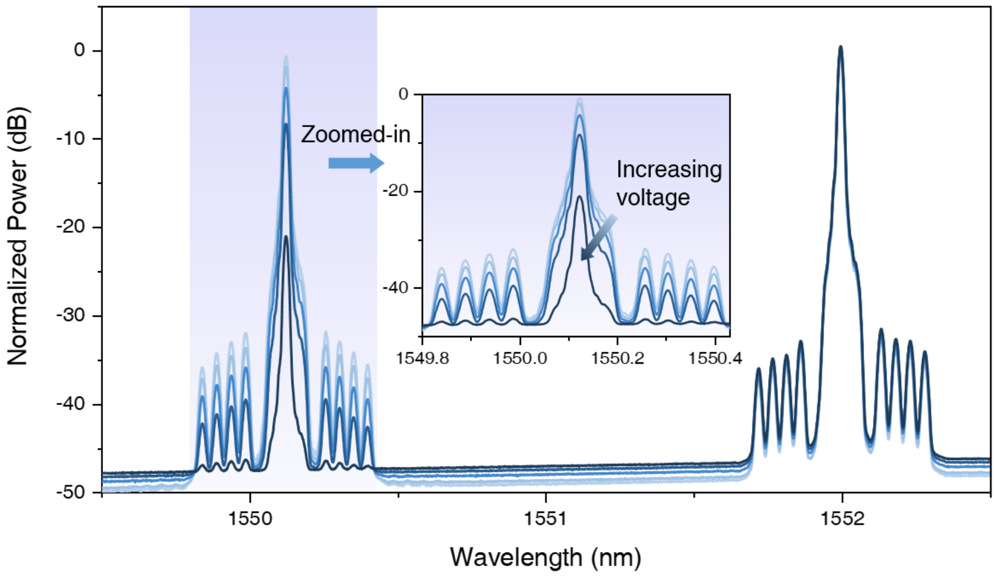

Fig. S14. Optical spectra of two wavelength channels (1550.12 nm and 1552 nm) after on-chip addition. The applied voltage on MUL-1 progressively attenuates all subcarriers on the 1550.12 nm carrier, demonstrating simultaneous weighting across the entire subcarrier set.

To further visualize the MAC process in the optical domain, we record the output spectra after the on-chip multipliers and adder under different programmed weights. Fig. S14 shows the combined optical spectrum after on-chip addition for two input optical carriers centered at 1550.12 nm and 1552 nm, each carrying four FDM subcarriers. The weighting functions are implemented by the on-chip multipliers MUL-1 and MUL-2. For clarity, only MUL-1 is tuned, while MUL-2 is kept unchanged. As the bias voltage applied to MUL-1 increases, the power of all subcarriers carried by the 1550.12 nm channel decreases simultaneously, whereas the 1552 nm channel remains essentially unchanged. This behavior confirms that the photonic multiplier imposes a nearly uniform weight across the full RF subcarrier band carried by a single optical channel.

C. Spectral parameter optimization for FDM-based computation

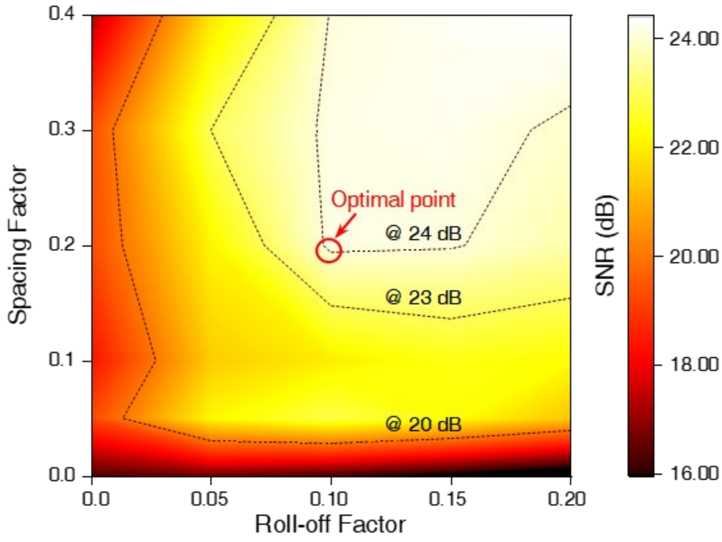

Fig. S15. Discussion of optimal roll-off and spacing factors for FDM-based computation.

In this section, we optimize the signal roll-off factor and the subcarrier spacing factor (SF) for FDM-based computation through 4-subcarrier optical addition experiments. Each subcarrier carries an 8-level high-speed signal at a symbol rate of 2 GBaud. Here, the subcarrier spacing is defined as (1+SF) × Baudrate. Fig. S15 shows the measured SNR map as a function of the signal roll-off factor and SF. It can be observed that small SF values result in spectral overlap between adjacent subcarriers, leading to degraded SNR. Conversely, large SF values cause unnecessary spectral gaps and reduce the overall spectral efficiency. The optimal performance is achieved at a roll-off factor of 0.1 and SF = 0.2, which provides a good balance between spectral efficiency and interference suppression.

D. RF-bandwidth scaling of FDM-based computation

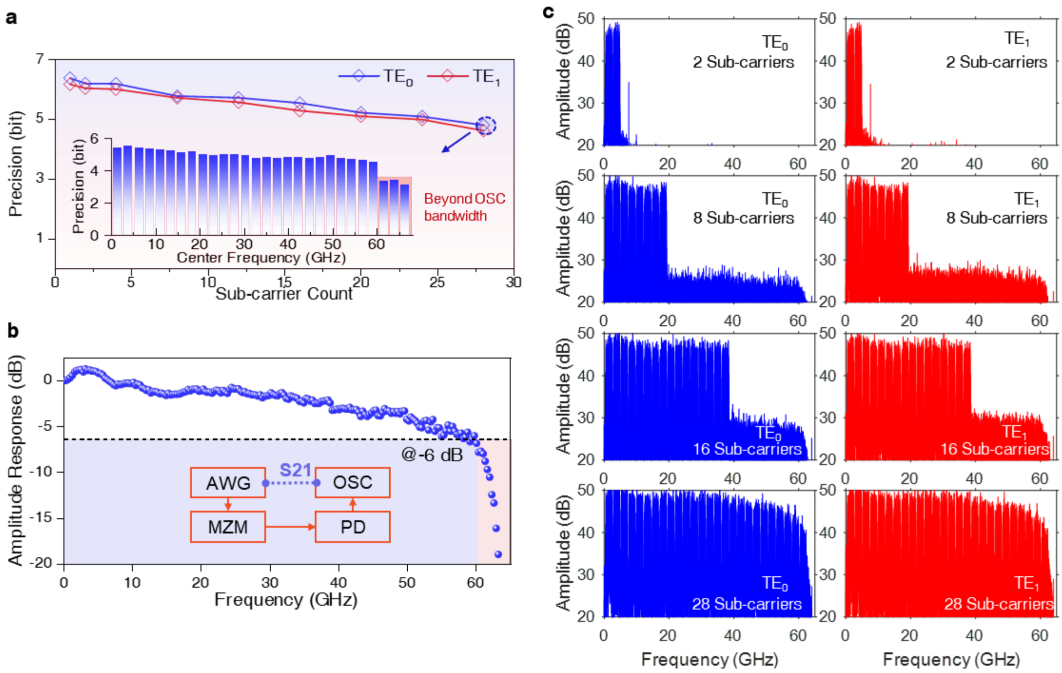

Fig. S16. Influence of occupied RF bandwidth on computational precision at a fixed per-subcarrier symbol rate of 2 GBaud. (a) Measured addition precision for TE_0 and TE_1 modes as the number of subcarriers increases from 1 to 28; inset: Precision distribution of the 28-subcarrier case for the TE_0 mode. **(b)** System end-to-end frequency response (S21). The red-shaded area denotes the high-attenuation region caused by the rapid roll-off of the oscilloscope's frequency response. **(c)** Measured RF spectra after PD detection for 2, 8, 16, 28-subcarrier cases, for TE_0 (blue) and TE_1 (red) modes, respectively.

To further clarify the scaling limit of RF parallelism, we investigate the impact of increasing the total occupied RF bandwidth on computational accuracy. In this measurement, the per-subcarrier bandwidth is kept fixed, while the number of subcarriers is gradually increased from 1 to 28 to expand the total RF bandwidth and evaluate the performance of FDM-based optical addition. Each subcarrier carries an 8-level signal at a symbol rate of 2 GBaud. The computational accuracy is quantified by the average precision over all subcarriers. The measured precision as a function of subcarrier number is shown in Fig. S16(a). As the spectrum extends toward higher frequencies, the signals experience stronger RF attenuation and increased high-frequency noise, resulting in an evident reduction in the average precision. In particular, the oscilloscope used in the experiment has an effective cutoff bandwidth of approximately 60 GHz, beyond which the frequency response rolls off rapidly, as shown in Fig. S16(b). Consequently, the outermost subcarriers beyond 60 GHz exhibit noticeably degraded precision, as highlighted in the inset of Fig. S16(a).

Fig. S16(c) further shows the measured spectra of the received signals after optical addition for different subcarrier counts. As the number of subcarriers increases, the occupied RF bandwidth expands progressively toward the high-frequency region, while the degradation becomes pronounced only when the spectrum approaches and exceeds the system bandwidth limit. Complementary to the fixed-bandwidth study in the main text, these results indicate that the present RF parallelism of the FieldCore is limited primarily by the available analog bandwidth of the experimental hardware, rather than by any fundamental limit on the number of FDM subcarriers, and could therefore be further scaled with higher-bandwidth optoelectronic components.

E. Supplementary results for the fixed-bandwidth FDM experiment

As a supplement to the fixed-bandwidth study in the main text, Fig. S17 provides the corresponding precision, spectral and waveform details. Fig. S17(a) summarizes the mean precision for the TE_0 and TE_1 modes as the number of subcarriers increases under a fixed total RF bandwidth of 60 GHz, together with the corresponding baudrate of each subcarrier. Fig. S17(b) shows the received RF spectra after PD detection for the 4-, 32-, 64-, and 100-subcarrier cases, confirming proper multiplexing for both modes within the occupied bandwidth. Fig. S17(c) further presents the recovered temporal waveforms for the 4-subcarrier case of the TE_1 mode, which agree well with the corresponding ideal results.

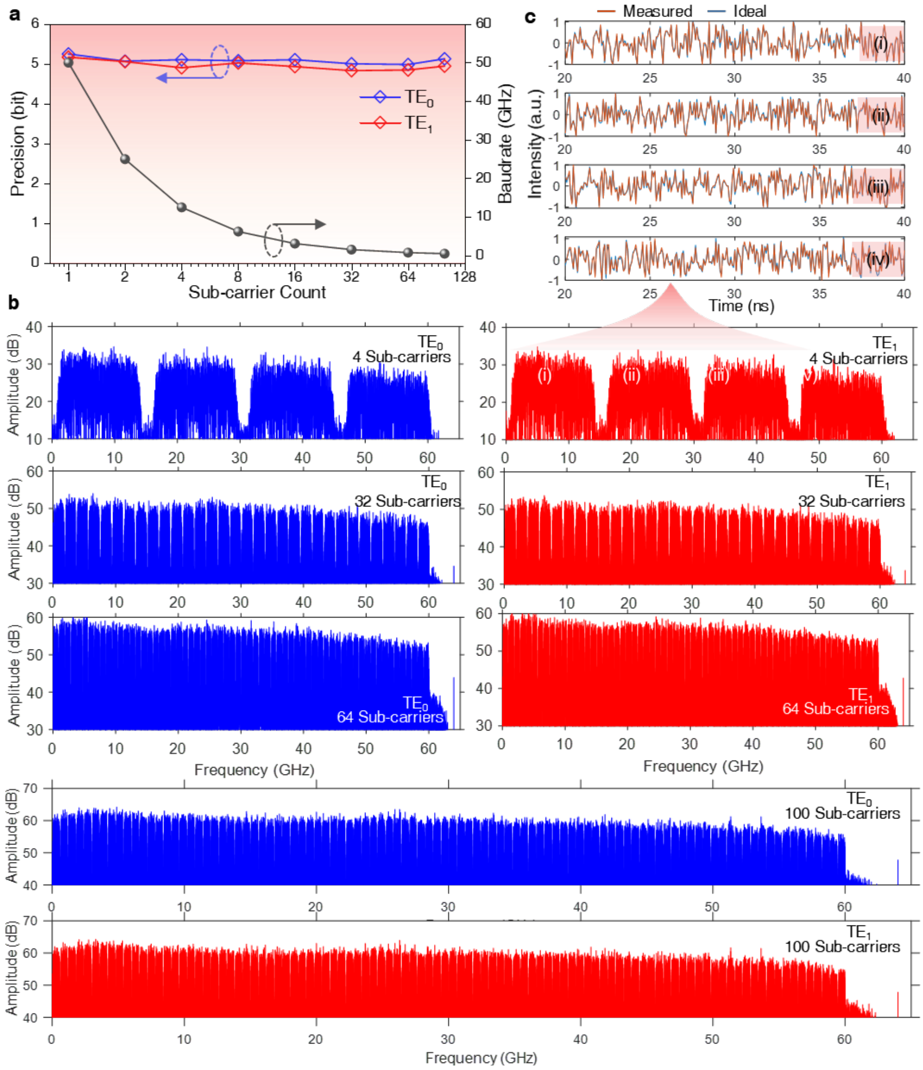

Fig. S17. Influence of subcarrier number under a fixed total RF bandwidth of 60 GHz. (a) Measured addition precision (left-axis) for TE_0 and TE_1 modes and the corresponding baudrate of each subcarrier (right-axis) as the number of subcarriers increases. **(b)** Received RF spectra for 4, 32, 64, and 100-subcarrier cases after PD detection for the two modes, respectively. **(c)** Measured temporal waveforms and the ideal waveforms for the 4-subcarrier case of the TE_1 mode.

Supplementary Note 5: Detailed Results of Arithmetic Operations Across Different Baudrates

A. Supplementary MAC results across different baudrates

As a representative illustration of the MAC process, we first present the input waveforms and the corresponding experimental results obtained at a symbol rate of 10 GBaud. The target operation follows the expression $X_1 \cdot Y_1 + X_2$, where two 8-level stepwise signals, X_1 and X_2, are generated using an AWG and modulated onto two distinct optical carriers, both propagating in the same waveguide mode. The first signal undergoes programmable optical weighting via the inverse-designed multiplier, effectively implementing the multiplication $X_1 \cdot Y_1$. The weighted signal and the second optical carrier are subsequently combined through an optical adder to yield the final MAC output. In this experiment, the two input wavelengths are set to 1550 nm and 1552 nm, respectively. Figs. S18(a)-(c) show the ideal temporal waveforms of X_1, $X_1 \cdot Y_1$, and X_2, respectively. Figs. S18(d)-(e) display the experimentally measured MAC outputs for TE_0 and TE_1 modes, along with the calculated ground truth. The experimental results exhibit good agreement with the ideal waveforms, confirming the effectiveness of our photonic MAC implementation.

As a supplement to the baudrate sweep summarized in the main text, Fig. S19 provides detailed MAC results for both TE_0 and TE_1 modes at baudrates from 10 to 120 GBaud. Each subplot compares the measured and expected values and includes the corresponding error histogram with a Gaussian fit. These results indicate that the FieldCore maintains mode-consistent and approximately linear MAC behavior up to 120 GBaud, with the residual degradation at high baudrates arising mainly from finite system bandwidth and accumulated noise[10].

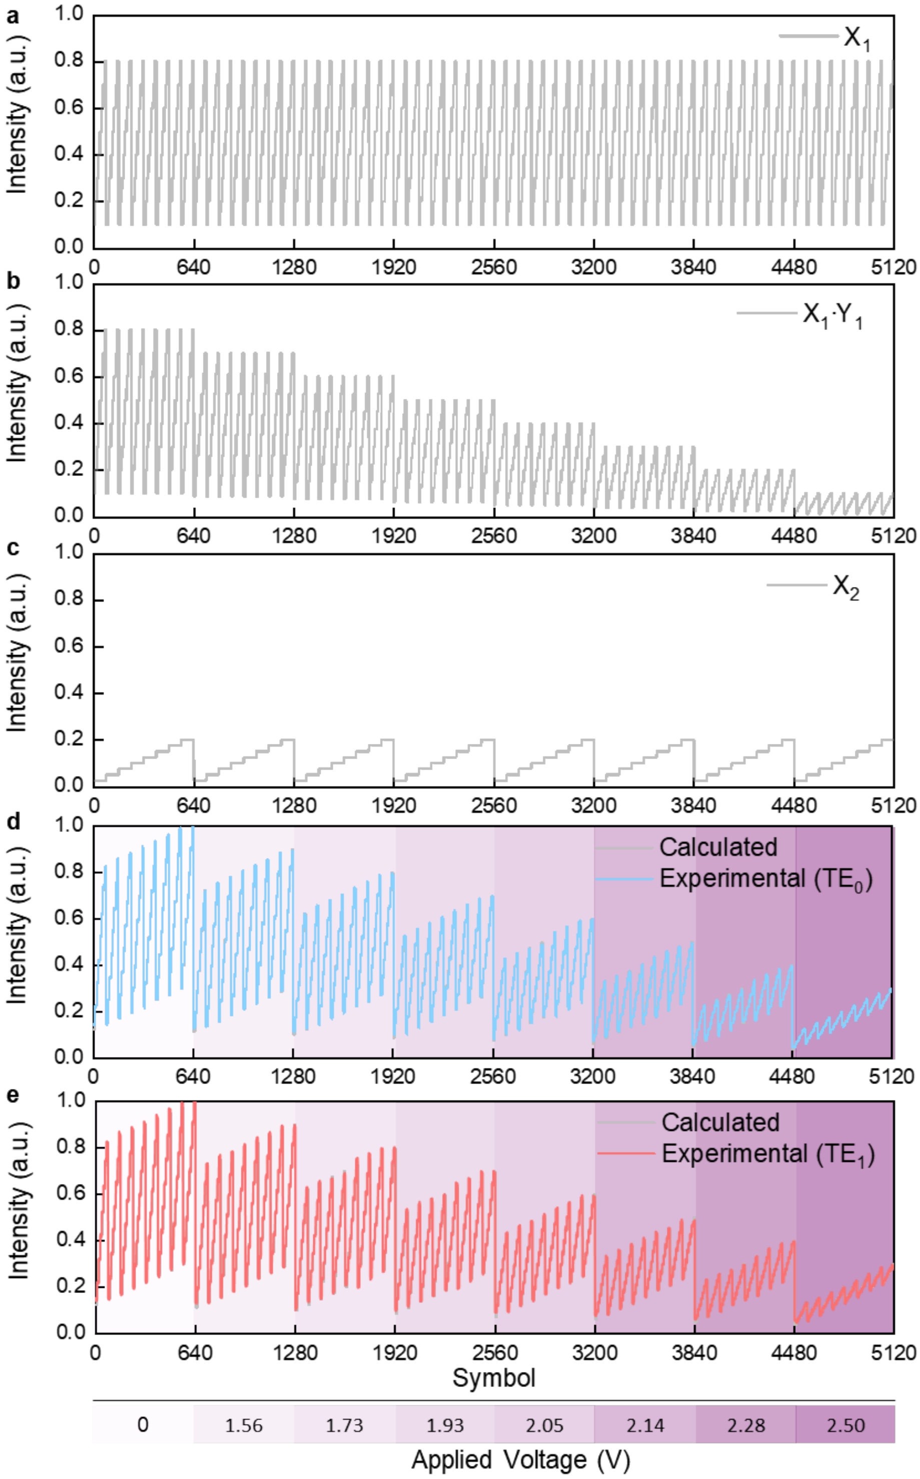

Fig. S18. Temporal waveforms illustrating the photonic MAC process at 10 GBaud. (a–c) Ideal waveforms of the two input operands and their intermediate result: **(a)** input X_1, **(b)** weighted signal $X_1 \cdot Y_1$, and **(c)** input X_2. **(d–e)** Experimentally measured MAC output waveforms for **(d)** TE_0 and **(e)** TE_1 modes, overlaid with calculated ground truth. The background color indicates the applied voltage on the multiplier, corresponding to different weight levels.

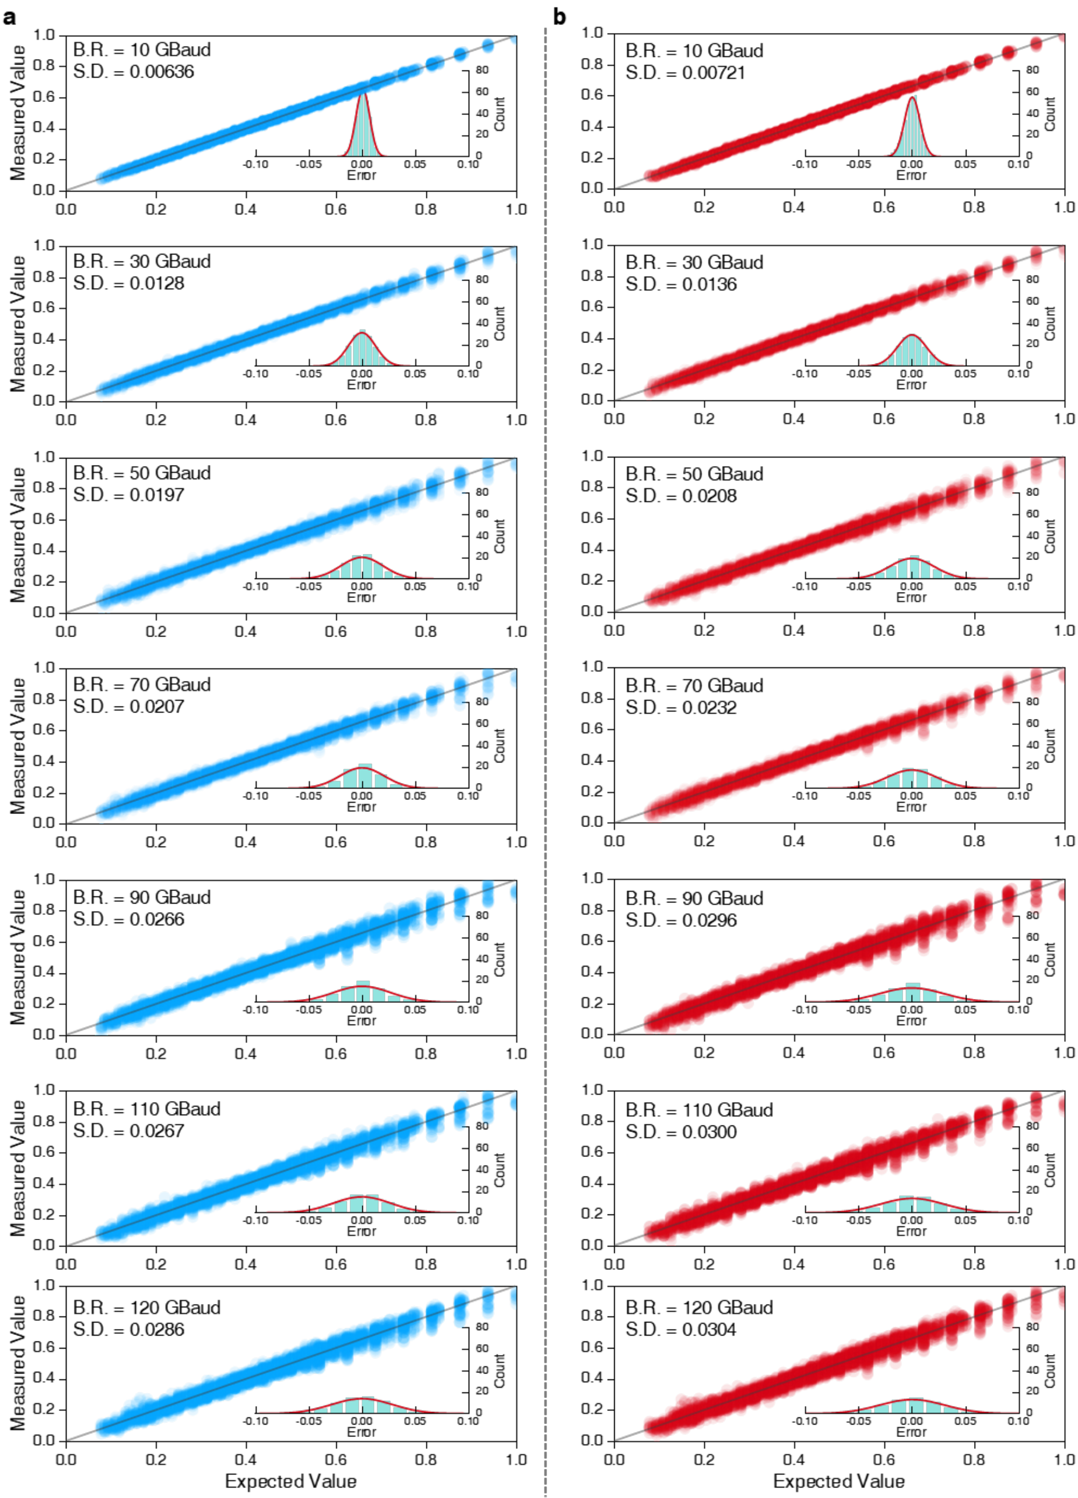

Fig. S19. Performance of photonic MAC operations at different baudrates. (a-b) Scatter plots of measured versus expected MAC results under **(a)** TE_0 and **(b)** TE_1 modes, with baudrates ranging from 10 GBaud to 120 GBaud. Insets show the corresponding error histograms and Gaussian fits. B.R., baudrate; S.D., standard deviation.

B. Measured bit-precision results under different baudrates

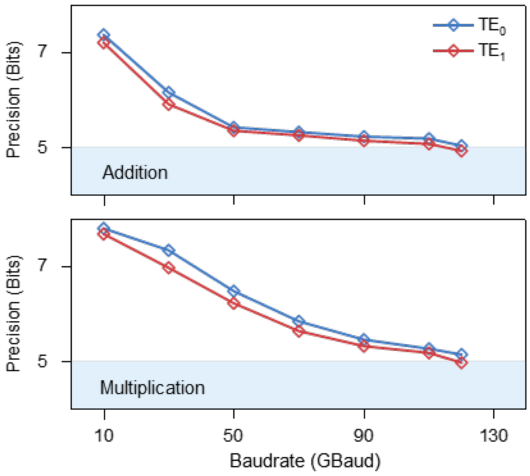

Fig. S20. Measured bit-precision for the two modes across different baudrates for addition and multiplication.

As a supplement to the baudrate-dependent MAC results presented in the main text, we further characterize the bit precision of optical addition and multiplication operations from 10 to 120 GBaud, as shown in Fig. S20.

Supplementary Note 6: Implementation details of FieldCore

A. Chip layout and packaging details

Fig. S21 shows the microscope image of the fabricated FieldCore chip together with enlarged views of representative functional regions. The overall photograph reveals the full device footprint, including the optical input/output coupling regions, the central computing array, and the surrounding electrical access area. Enlarged views highlight several representative building blocks, including the mode MUX, multiplier, and waveguide crossing. Fig. S22 further illustrates the electrical interfacing design used for external bias access. The left panel shows the chip layout with the electrical pad arrays distributed along the upper and lower sides of the circuit. The middle panel extracts the pad layout alone and labels the pad numbers for channel assignment. The right panel shows the routing design that connects the chip pads to the probe-pin interface, enabling practical access from the external multi-channel DC control system to the densely distributed electrical ports on the FieldCore.

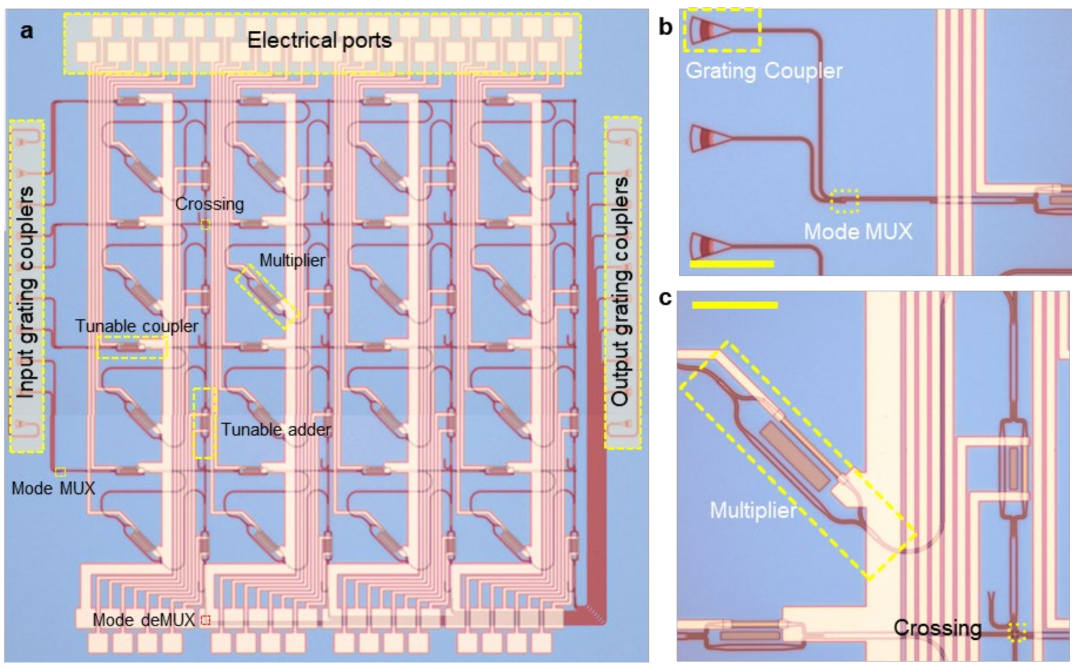

Fig. S21. Microscope images of the fabricated FieldCore chip. (a) The overall view; **(b)** The zoomed-in view of the input area; **(c)** The zoomed-in view of the computing area.

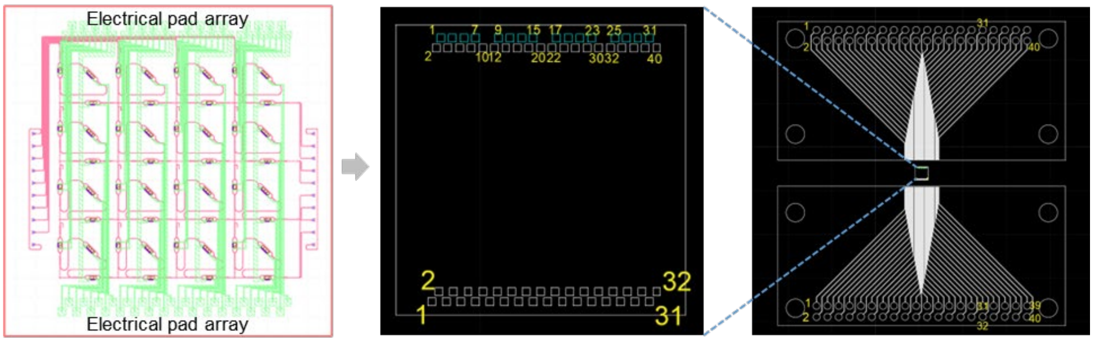

Fig. S22. Photonic circuit layout and corresponding PCB interfacing design.

B. Experimental configuration for electrical biasing

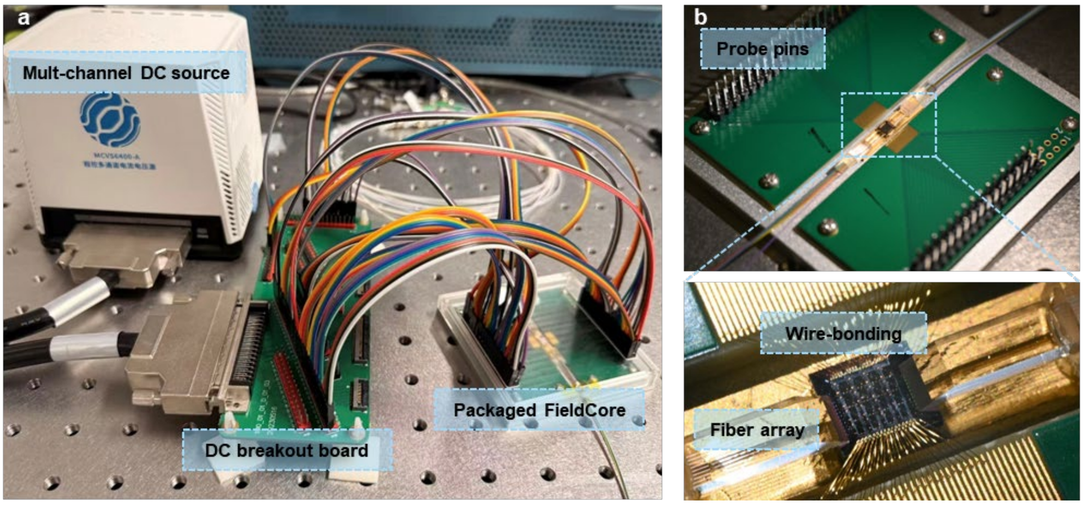

Fig. S23. Electrical biasing and packaging configuration of the FieldCore. (a) Photograph of the electrical biasing setup. **(b)** Close-up views of the PCB board and the packaged chip.

To provide electrical biasing for the FieldCore (Fig. S23), a multi-channel DC source is connected to a DC breakout board, and the distributed bias lines are routed to the packaged device via jumper wires. The electrical connection to the chip is realized through probe pins and wire-bonded pads, enabling independent control of the on-chip tunable elements.

C. Architecture of the photonic MAC unit for arithmetic validation

To quantitatively test the performance of high-speed arithmetic operations, we also design and implement an MAC unit that supports addition, multiplication, and MAC operations. The fabricated chip, as shown in Fig. S24, integrates an array of grating couplers, a pair of mode multiplexers, a mode demultiplexer, two optical multipliers, and an optical adder based on a multimode interference (MMI) coupler. The chip is packaged using wire bonding for electrical interfacing and fiber array coupling for optical I/O.

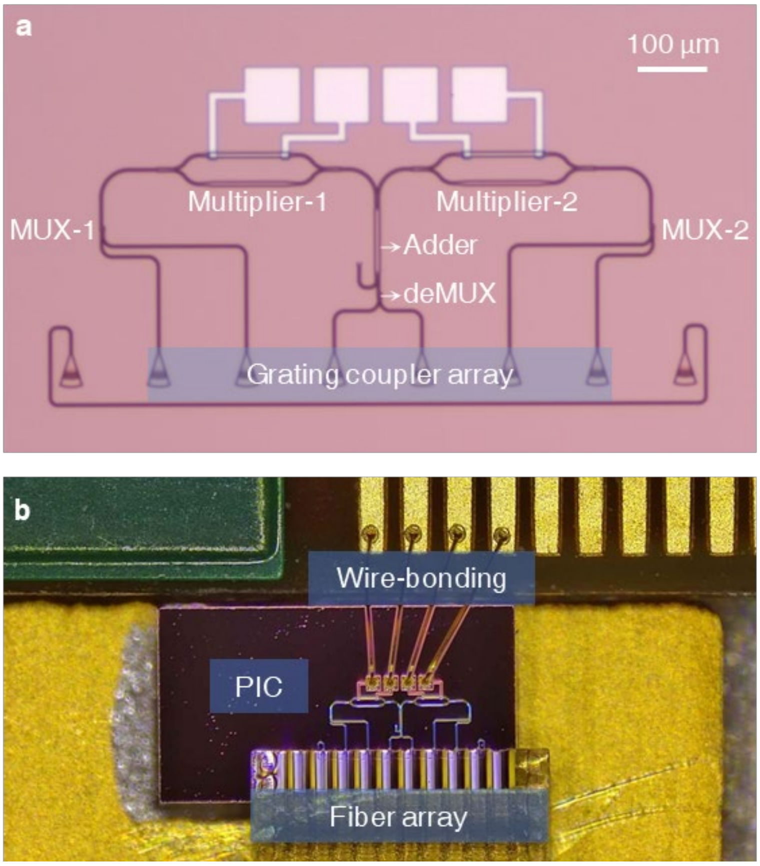

Fig. S24. Fabricated multidimensional photonic MAC core. (a) Optical microscope image of the MAC core chip, consisting of a grating coupler array, two mode multiplexers, two optical multipliers, a mode demultiplexer, and a dual-mode MMI as an optical adder. **(b)** Photograph of the packaged chip with wire-bonded electrical contacts and fiber array coupling for optical input/output.

Supplementary Note 7: Demonstration of Image Convolution

A. Grayscale image convolution across different baudrates

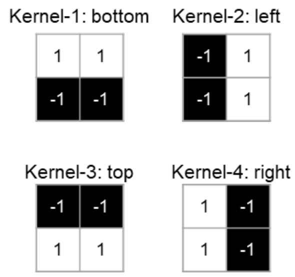

Fig. S25. Convolution kernels used in the grayscale convolution experiment.

Four representative 2 × 2 convolution kernels are used in the grayscale image-processing experiment, corresponding to bottom, left, top and right directional gradients, as shown in Fig. S25. The measured convolution performance for the four kernels at different baudrates is summarized in Fig. S26, which provides the peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM) values for the kernel-specific outputs under both modes[11]. Representative convolution results at 10 GBaud for the four directional kernels are shown in Fig. S27.

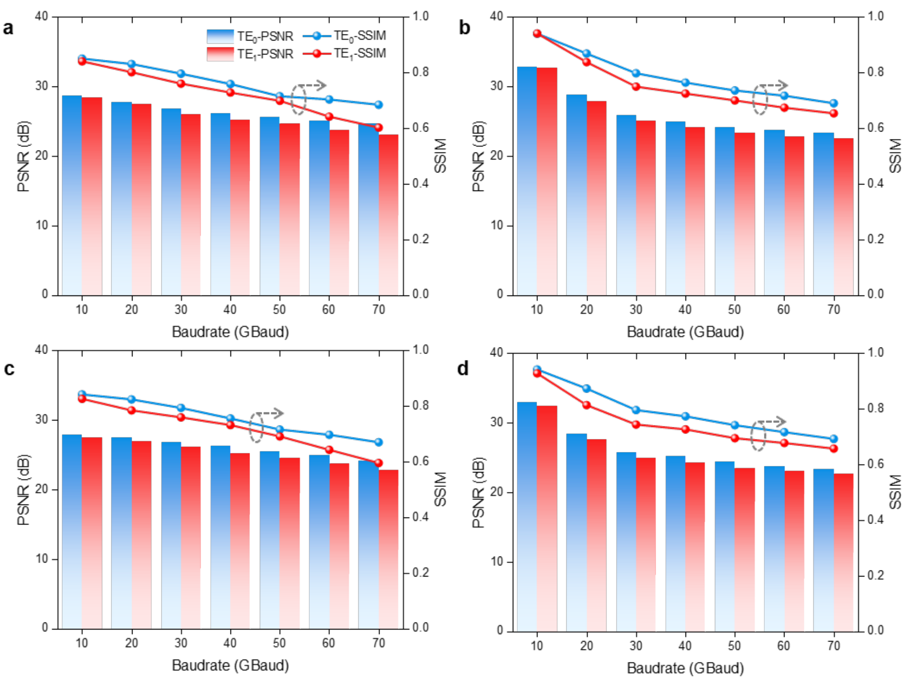

Fig. S26. Measured convolution performance across baudrates for (a) kernel-1, **(b)** kernel-2, **(c)** kernel-3, and **(d)** kernel-4 for the TE_0 and TE_1 modes. The bar plots represent the PSNR values, while the line plots represent the corresponding SSIM values.

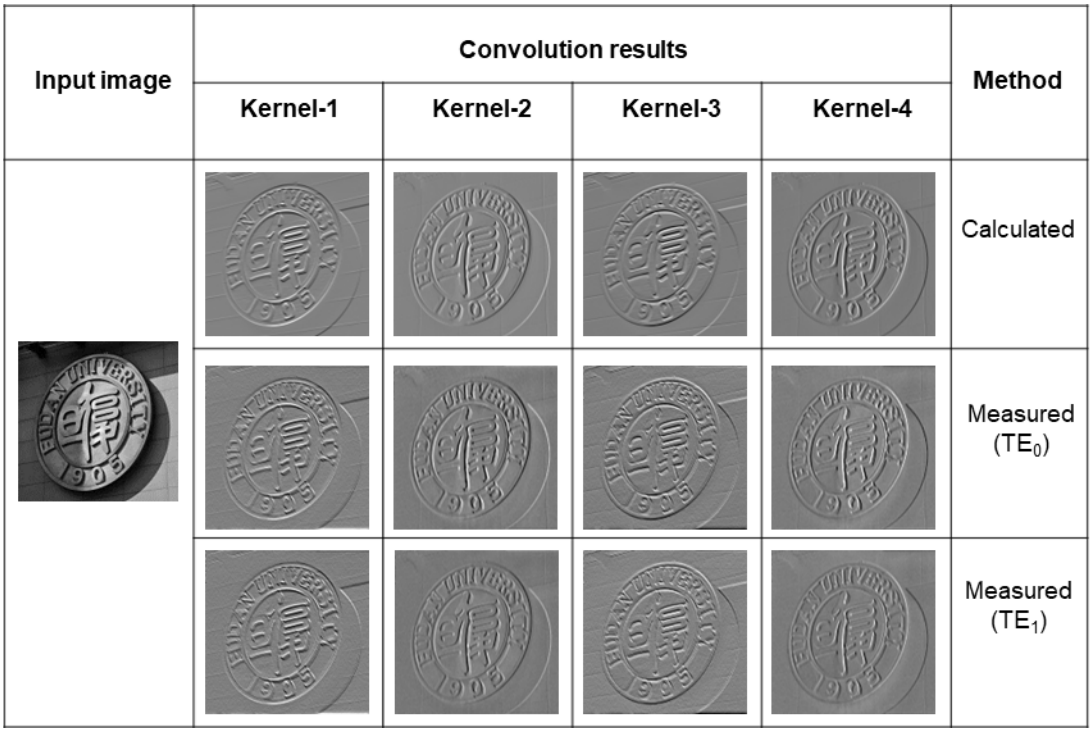

Fig. S27. Grayscale convolution outputs for the four kernels at 10 GBaud.

After obtaining the convolution outputs of the four directional kernels, we further combine them to generate a two-dimensional edge map using the standard gradient-magnitude formulation, analogous to Prewitt-type edge detection. Specifically, the four convolution outputs associated with the bottom, left, top, and right kernels are denoted as Y_{bottom}, Y_{left}, Y_{top}, and Y_{right}, respectively. The vertical and horizontal gradients are then computed as $G_y = \frac{Y_{bottom} - Y_{top}}{2}, G_x = \frac{Y_{right} - Y_{left}}{2}$, for both the ideal digital reference and the measured optical results. The final edge map is then obtained from the gradient magnitude, $E = \sqrt{G_x^2 + G_y^2}$, followed by intensity normalization. Fig. S28 shows the combined edge maps at processing rates of 10, 30, and 50 GBaud, confirming the feasibility of performing parallel, high-speed optical image processing using the FieldCore.

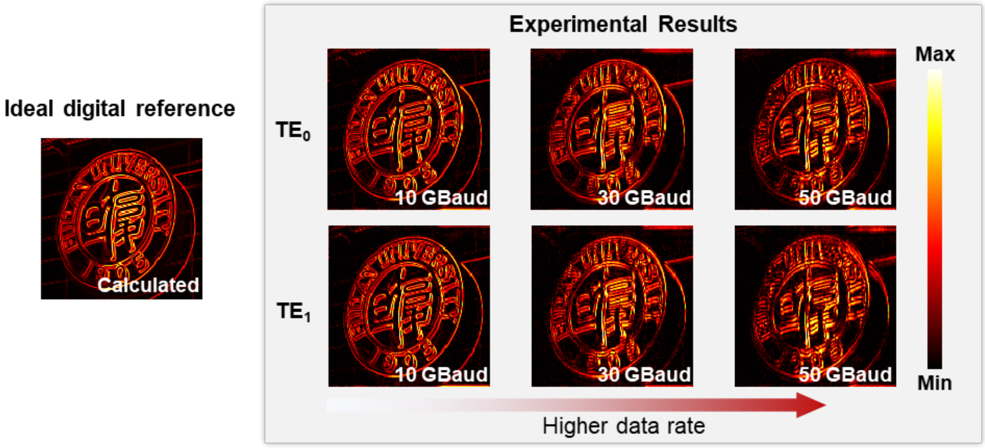

Fig. S28. Combined edge maps from the four convolution outputs. Comparison between the calculated edge map and the experimentally measured edge maps obtained at baudrates of 10, 30, and 50 GBaud for the TE_0 and TE_1 modes.

B. Supplementary RGB image convolution results

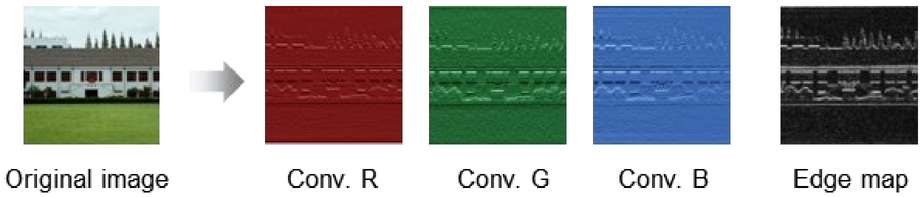

Fig. S29. The original RGB images and the corresponding convolution results for TE_1 mode.

For clarity, Fig. S29 presents the corresponding RGB convolution results and reconstructed edge map for the TE_1 mode. The TE_1 results closely match those obtained in the TE_0 mode shown in the main text, further confirming mode-consistent convolution performance under FDM operation.

Supplementary Note 8: Demonstration of Parallel AI inference

A. Experimental setup for fully multiplexed photonic computation

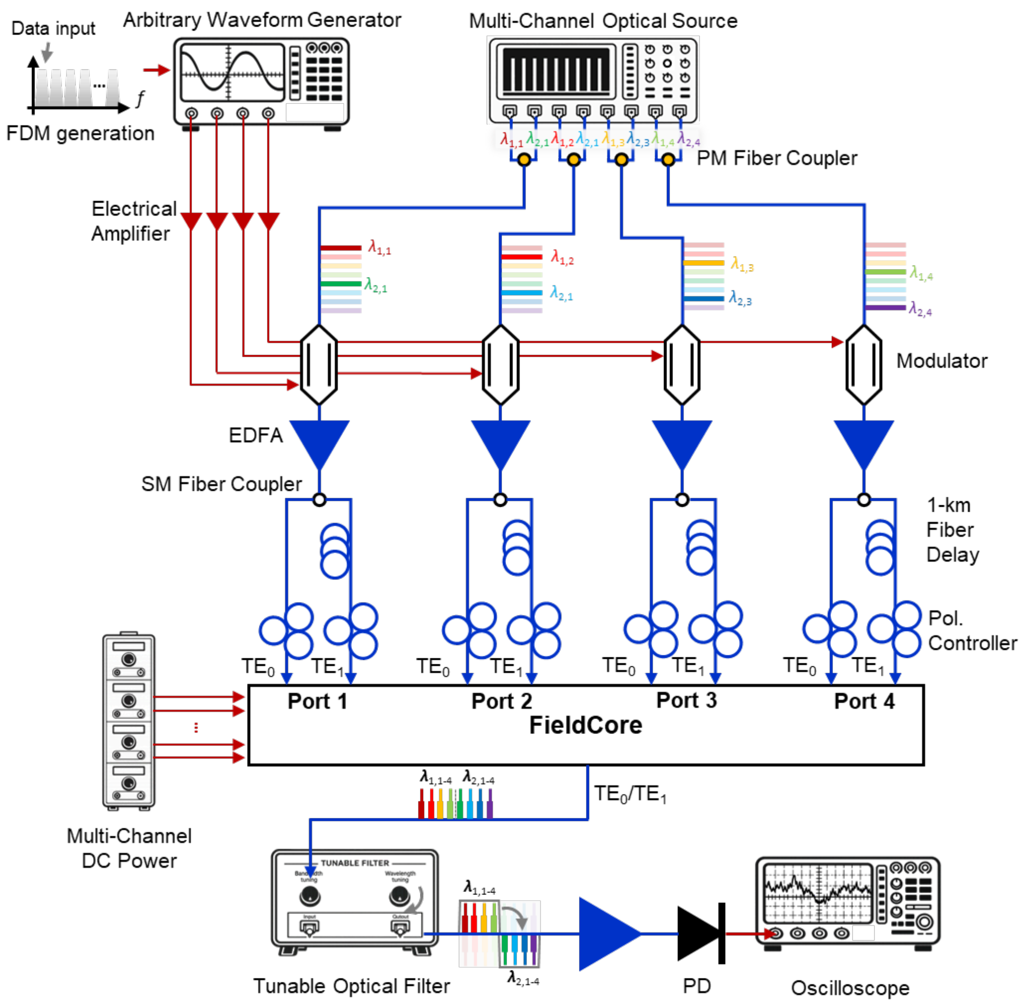

Fig. S30. Experimental setup for multi-dimensional photonic computing. FDM signals are electrically generated, loaded onto WDM carriers, and launched into decorrelated TE_0 and TE_1 mode channels. Blue lines denote optical paths, and red lines denote electrical paths.

Fig. S30 illustrates the experimental setup for fully multiplexed multi-dimensional computation. In the wavelength dimension, eight optical carriers are arranged into two groups, $\{\lambda_{1,1}, \lambda_{1,2}, \lambda_{1,3}, \lambda_{1,4}\}$ and $\{\lambda_{2,1}, \lambda_{2,2}, \lambda_{2,3}, \lambda_{2,4}\}$. Each four-wavelength group corresponds to one four-element spatial input vector, with the four wavelengths assigned to Ports 1-4, respectively. This grouping is adopted so that optical accumulation across the four spatial ports remains free from phase-sensitive coherent interference during direct detection[8]. For each spatial port r, the two carriers $\lambda_{1,r}$ and $\lambda_{2,r}$ are combined and modulated together, generating a two-channel WDM input.

To establish the mode dimension, each modulated optical stream is divided into two branches, one of which passed through a 1-km fiber delay before injection. For a laser with Lorentzian linewidth $\Delta\nu$, the coherence length can be approximated as[12]

$$L_c = \frac{c}{\pi n_g \Delta\nu}, \tag{S8.1}$$

where c is the speed of light in vacuum, n_g is the group index of the fiber, and $\Delta\nu$ is the laser linewidth. For the 100-kHz-linewidth laser used in the experiment, the corresponding coherence length is approximately 0.65 km. The applied delay (1 km) therefore enables decorrelated dual-mode loading for the TE_0 and TE_1 channels.

At the receiver, a bandwidth- and wavelength-tunable optical filter is used for channel readout. The filter bandwidth is adjusted to cover four wavelength channels, such that one complete wavelength group could be selected at a time. The passband is first tuned to select $\{\lambda_{1,1}, \lambda_{1,2}, \lambda_{1,3}, \lambda_{1,4}\}$, and then shifted to $\{\lambda_{2,1}, \lambda_{2,2}, \lambda_{2,3}, \lambda_{2,4}\}$. Combined with manual kernel selection and separate detection of the two guided-mode channels after on-chip mode demultiplexing, this readout scheme enabled channel-resolved measurement of different channels without changing the simultaneous on-chip computation. Specifically, eight optical carriers at 1547, 1548, 1549, 1550, 1551, 1552, 1553, and 1554 nm were used in the following experiments. As shown in Fig. S31, the top panel presents the unfiltered spectrum containing all channels, while the lower two panels show the spectra after selecting wavelength group 1 (λ_1) and wavelength group 2 (λ_2), respectively, corresponding to 1547-1550 nm and 1551-1554 nm.

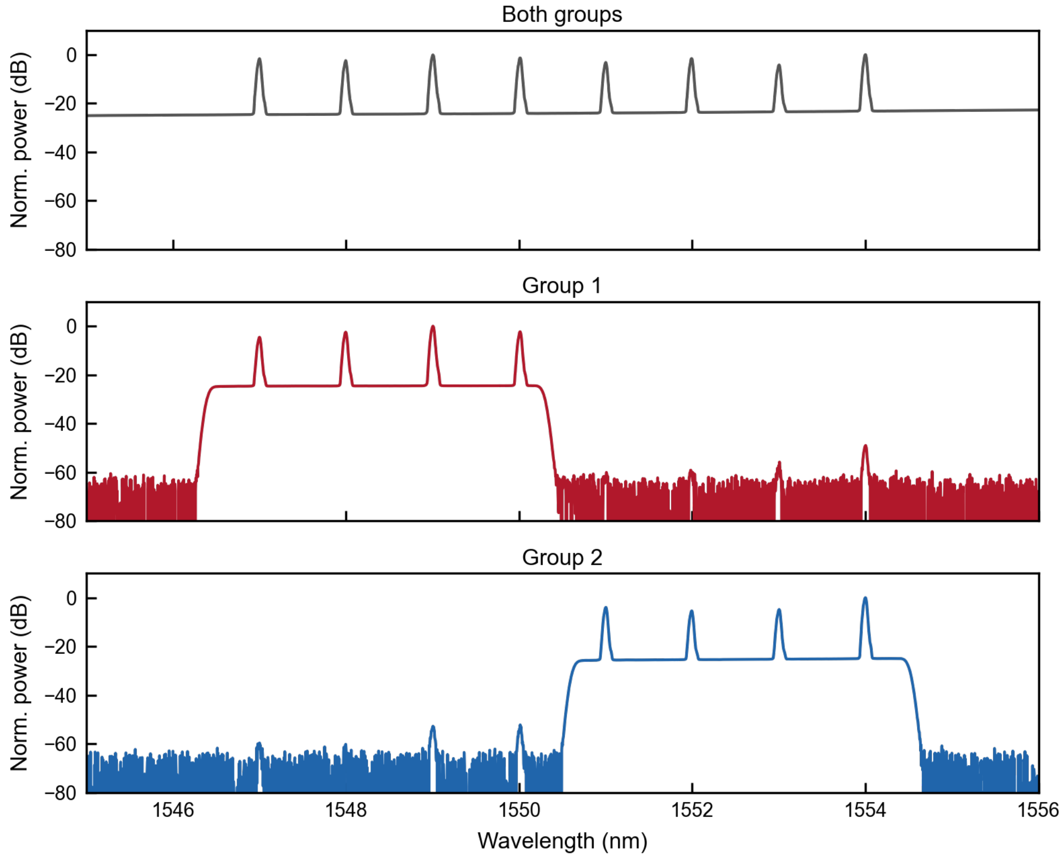

Fig. S31. Measured optical spectra before and after wavelength-group selection at the receiver. The top panel shows the unfiltered spectrum containing all eight carriers, while the lower panels show wavelength group 1 and group 2, corresponding to 1547–1550 nm and 1551–1554 nm, respectively.

B. Handwritten-digit recognition

This note provides additional experimental details for the handwritten-digit recognition results shown in Fig. 4d, with emphasis on the data configuration and the channel-resolved fidelity of the recovered convolution features. The input data consist of 28 × 28 grayscale handwritten digits. Each image is processed by four fixed 2 × 2 kernels corresponding to bottom-edge, left-edge, top-edge, and right-edge detection. Valid convolution followed by ReLU produces four 27 × 27 feature maps, which are concatenated and flattened into a 2,916-dimensional feature vector for classification. For the digital baseline, training is performed on the ideal digitally computed features using an 80%/10%/10% training/validation/test split. The classifier is optimized with Adam for 80 epochs using a learning rate of 1×10^{-3}, an L2 regularization factor of 1 ×

10^{-4}, a batch size of 128, and shuffling at every epoch.

During the experiment, two wavelength groups (1547-1550 nm, and 1551-1554 nm), two guided modes (TE_0 and TE_1), and 50 RF subcarriers are jointly employed, with each RF subcarrier carrying one handwritten-digit image. In the RF dimension, each subcarrier operates at 0.25 GBaud, with a spacing of 0.25 GHz and a starting frequency of 1 GHz, yielding 50 subcarriers spanning from 1 to 13.25 GHz. Fig. S32 shows the electrical spectrum of the received FDM signal after convolution, illustrating the 50-subcarrier RF multiplexing configuration used in the handwritten-digit recognition experiment. The evenly spaced subcarriers agree with the designed frequency allocation and confirm successful simultaneous loading and readout across the full RF band. The non-flat, spike-like spectral profiles arise because the transmitted image signals contain substantial DC and low-frequency components.

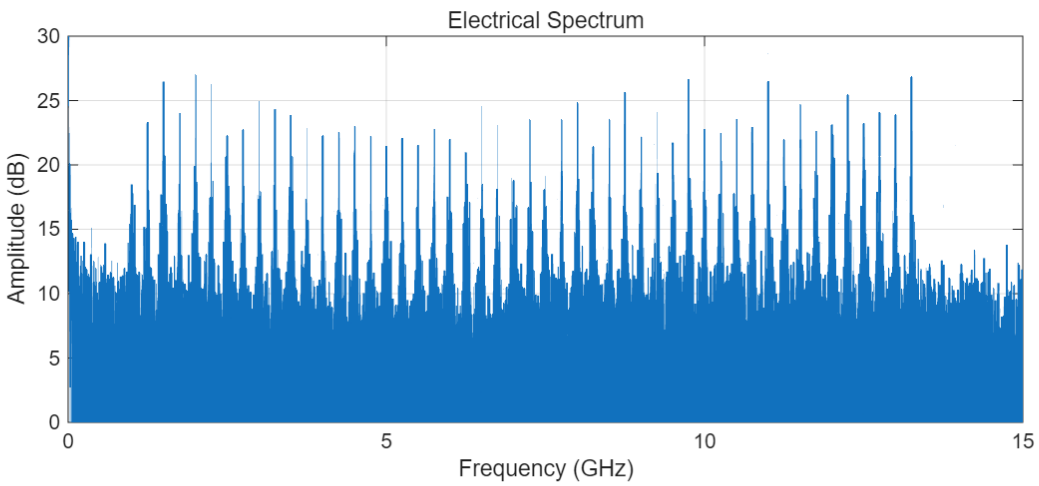

Fig. S32. Measured electrical spectrum of the received FDM signal in the handwritten-digit recognition experiment. Fifty RF subcarriers are simultaneously loaded, with a baudrate of 0.25 GBaud per subcarrier, a spacing of 0.25 GHz, and a starting frequency of 1 GHz.

Representative recovered convolution feature maps are shown in Fig. S33 for TE_0-λ_1 channel at RF-1, RF-25, and RF-50, corresponding to the low-, intermediate-, and high-frequency regions of the multiplexed RF band. For each selected RF subcarrier, the recovered feature maps of digits 0-9 under the four kernels are displayed. We further compare the experimentally recovered signals with the corresponding ideal target waveforms for a representative case. Fig. S34 shows the results for RF-25 under kernel

K1 in the TE_0-λ_1 channel, covering digits 0-9. The recovered waveforms closely follow the target traces for all digits, indicating that the temporal profiles of the convolution outputs are well preserved after fully multiplexed photonic processing.

To quantify the fidelity of the recovered convolution outputs, the PSNR is calculated with respect to the corresponding ideal digitally computed feature maps. Fig. S35 summarizes the channel-resolved PSNR values for all ten digits and all four kernels across the four wavelength-mode channel combinations. The error bars denote the standard deviation over the 50 RF subcarriers. The results confirm stable convolution fidelity across different digits, kernels, wavelength groups, guided modes, and RF subcarriers.

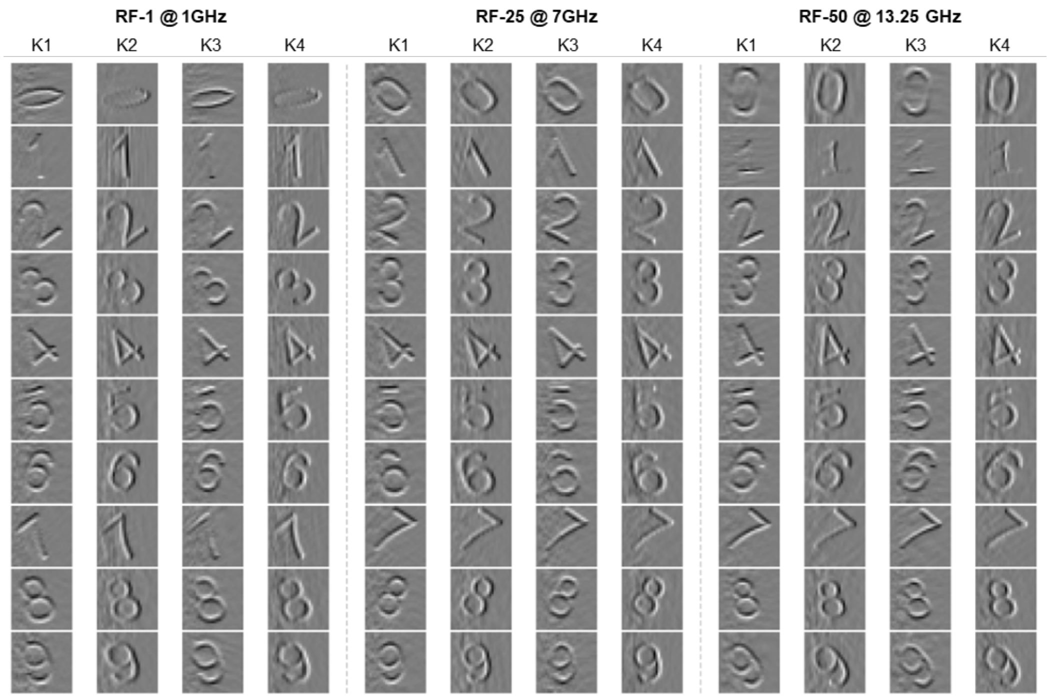

Fig. S33. Representative recovered convolution feature maps for handwritten-digit images in the TE_0-λ_1 channel at different RF subcarriers. Results are shown for RF-1, RF-25, and RF-50, corresponding to 1, 7, and 13.25 GHz, respectively. For each RF subcarrier, the recovered feature maps of digits 0–9 under the four directional kernels are displayed.

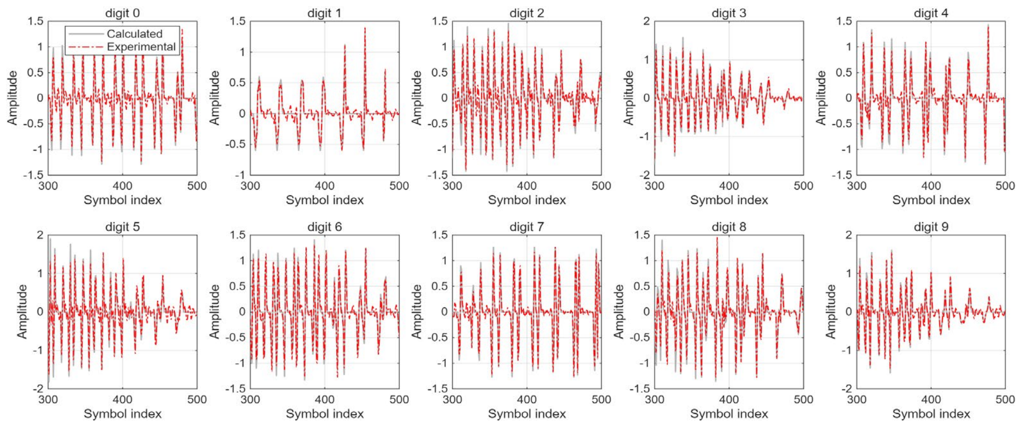

Fig. S34. Comparison between experimentally recovered and ideal target temporal waveforms of the convolution outputs. Symbol indices 300-500 are shown, as this segment contains more pronounced waveform features associated with the digit patterns.

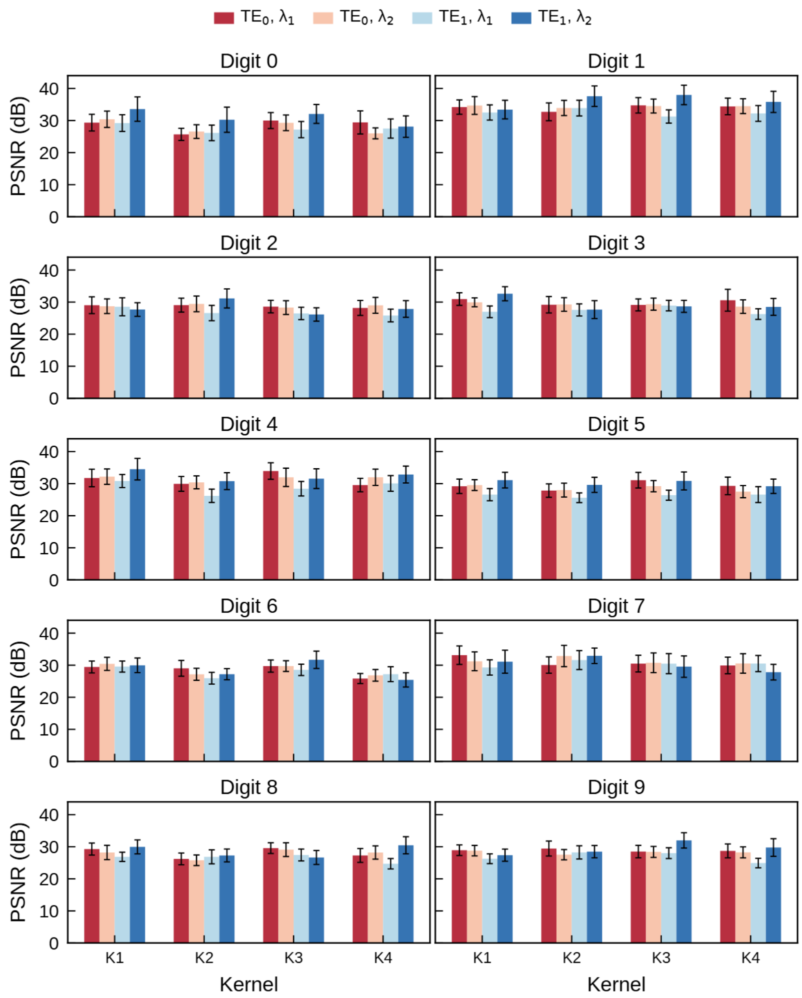

Fig. S35. Channel-resolved PSNR of recovered convolution feature maps in the handwritten-digit recognition experiment. Error bars denote the standard deviation over the 50 RF subcarriers. K1-K4 are four directional kernels corresponding to bottom, left, top, and right edge detection, respectively.

C. Hyperspectral image classification

The Indian Pines dataset comprises 145 × 145 pixels acquired by the AVIRIS sensor across 220 spectral bands spanning 400-2500 nm[13]. After removing 20 bands covering water absorption regions, the corrected dataset retains 200 bands. We select the 12 land-cover classes with sufficient sample counts, discarding four classes with fewer than 100 labelled pixels to ensure statistically reliable evaluation. Fig. S36 shows representative spectral signatures of the 12 classes, where the grey shaded regions indicate the excluded water absorption bands. For each labelled pixel, a 7 × 7 spatial neighborhood is extracted with mirror padding, yielding a 200 × 7 × 7 spectral-spatial input tensor. The 200 input channels are mapped one-to-one onto the 200 multiplexed optical channels of FieldCore, formed by 2 wavelength groups × 2 modes × 50 RF subcarriers. The schematic of the classification pipeline is shown in Fig. S37. Four shared 2 × 2 convolutional kernels are then applied to all 200 channels, generating an 800-channel feature map. Global average pooling followed by two fully connected layers with 128 and 12 neurons completes the classification pipeline. Training uses a per-class stratified split of 15% for training and 85% for testing.

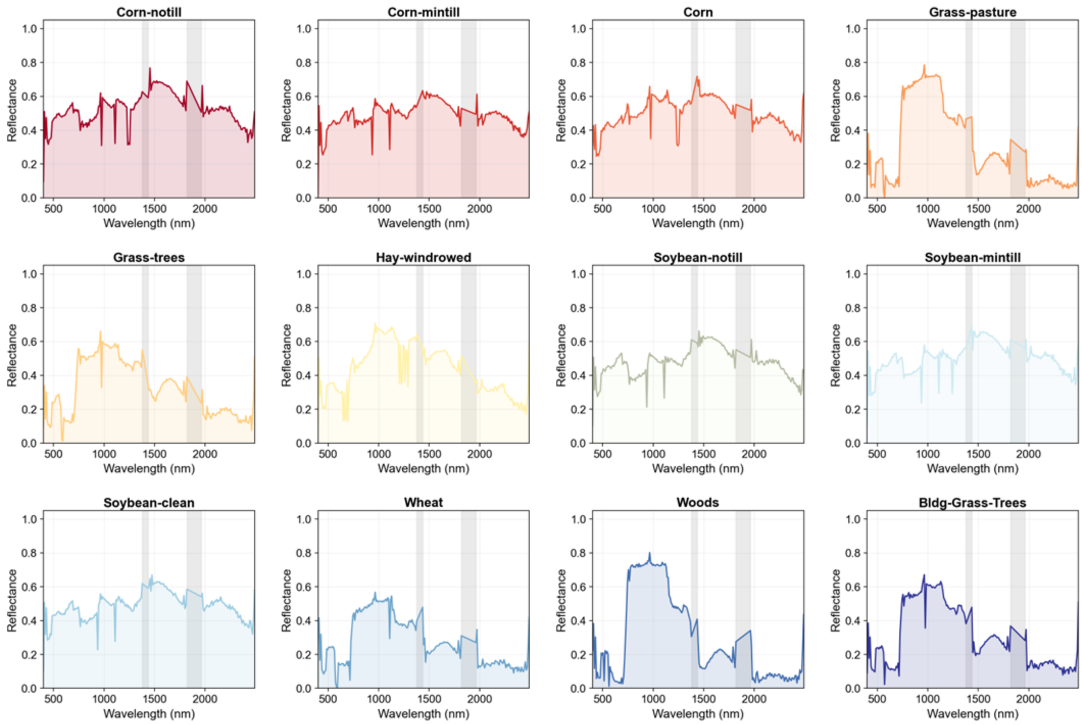

Fig. S36. Representative spectral signatures of the 12 land-cover classes in the Indian Pines dataset. Each curve shows the reflectance spectrum of a representative pixel from its class across 200 spectral bands (400–2500 nm). Grey shaded regions indicate water absorption bands removed in the corrected dataset.

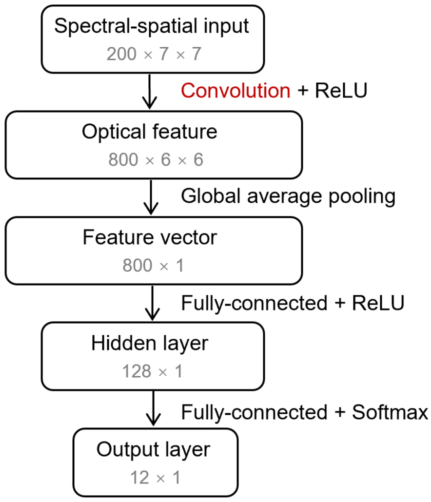

Fig. S37. The model architecture schematic for hyperspectral image classification.

To improve robustness against hardware noise, we employ noise-aware training (NAT)[14], in which Gaussian noise is injected at the output of the optical convolution layer during training, with the noise level scaled to the target effective precision N as $\sigma_{\mathrm{rel}} = 1/2^N$. For the digital precision sweep, a dedicated NAT model is trained and evaluated at each target effective precision. Fig. S38 shows the classification maps obtained under effective precisions ranging from 8-bit to 2-bit, together with the ground truth. The corresponding per-class confusion matrices at each precision level are shown in Fig. S39.

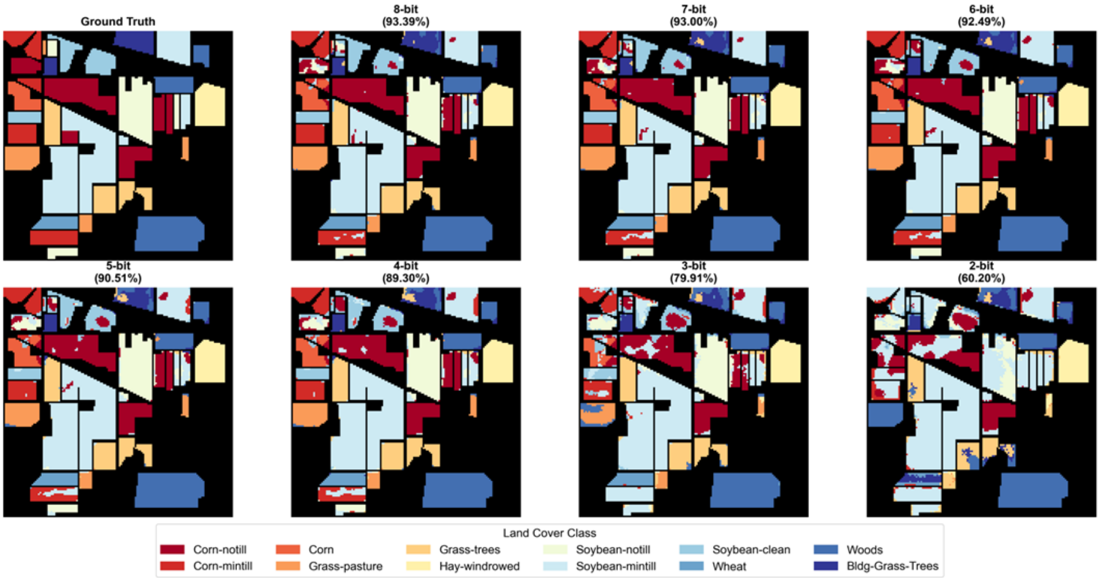

Fig. S38. Classification maps of the Indian Pines dataset under different effective precisions.

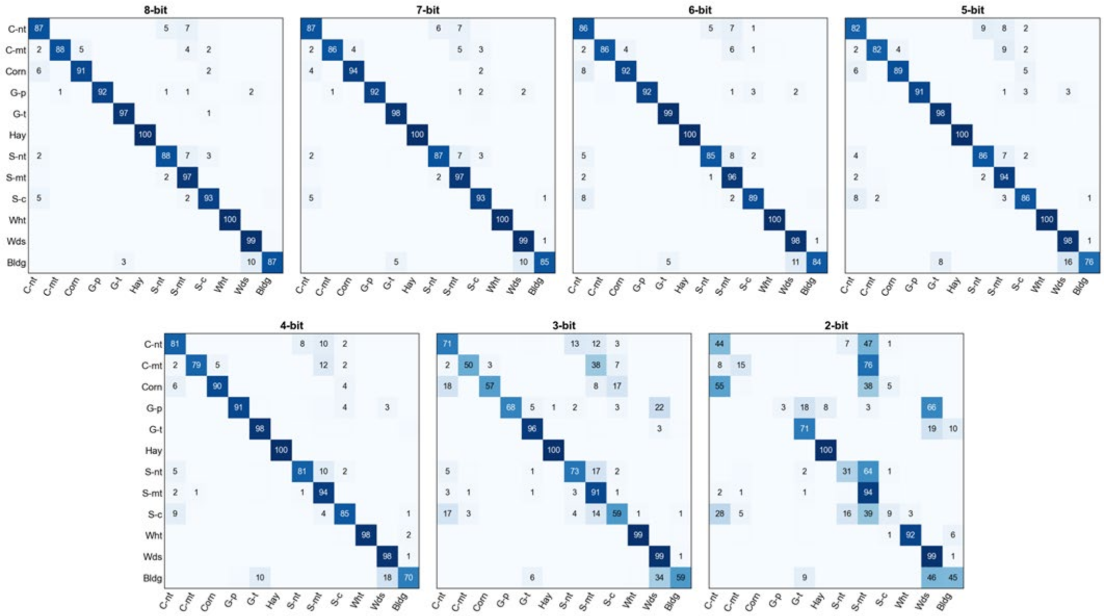

Fig. S39. Confusion matrices for the digital precision sweep. Class abbreviations: C-nt, Corn-notill; C-mt, Corn-mintill; Corn; G-p, Grass-pasture; G-t, Grass-trees; Hay, Hay-windrowed; S-nt, Soybean-notill; S-mt, Soybean-mintill; S-c, Soybean-clean; Wht, Wheat; Wds, Woods; Bldg, Buildings-Grass-Trees.

D. Mechanical fault diagnosis

We further evaluate FieldCore on the Case Western Reserve University (CWRU) bearing dataset[15], a widely used benchmark for vibration-based mechanical fault diagnosis. Vibration signals are acquired from a drive-end accelerometer at 12 kHz and categorized into 7 classes: one normal operating condition and six fault types comprising ball, inner race, and outer race defects at fault diameters of 0.007″ and 0.014″. Fig. S40 shows representative time-domain waveforms and their frequency-domain spectra for each class. Each raw signal is segmented into windows of 1024 samples with 50% overlap and independently normalized to [0, 1], yielding a 1-D input of dimension 1 × 1024. The dataset is split into 1,325 training and 568 test segments (70%/30%). Here 200 multiplexed channels of FieldCore serve a different role from the hyperspectral task: rather than carrying spectral bands of a single input, they enable instance parallelism: 200 vibration segments are loaded simultaneously across 2 wavelength groups × 2 guided modes × 50 RF subcarriers for batch-parallel inference in a single pass. During FieldCore inference, 200 random test segments are simultaneously loaded onto the 200 optical channels in each pass, and this process is repeated over 20 passes to accumulate statistically robust accuracy results.

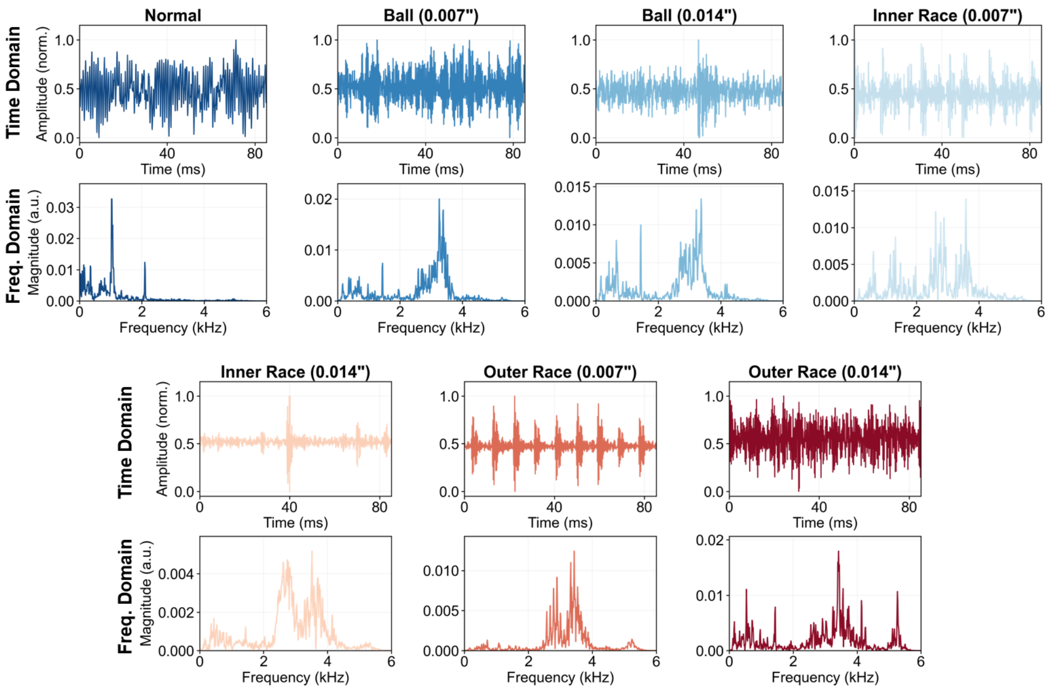

Fig. S40. Representative time-domain waveforms and frequency-domain spectra of the 7 fault classes.

The classification network (Fig. S41) consists of three 1-D convolutional layers, followed by global average pooling and a fully connected output layer. The first convolutional layer is implemented by FieldCore, which applies four kernels of size 4 shared across all 200 parallel instances, producing a 4 × 1021 feature map per segment. After ReLU activation and max pooling, the feature dimension is reduced to 4 × 255. Two subsequent convolutional layers then expand the feature channels from 4 to 16 and from 16 to 32, yielding feature maps of size 16 × 63 and 32 × 63, respectively. Global average pooling compresses the final feature map into a 32 × 1 feature vector, which is then fed to the output layer for seven-class classification. As in the hyperspectral task, NAT is applied with additive Gaussian noise injected at the output of the first convolutional layer during training.

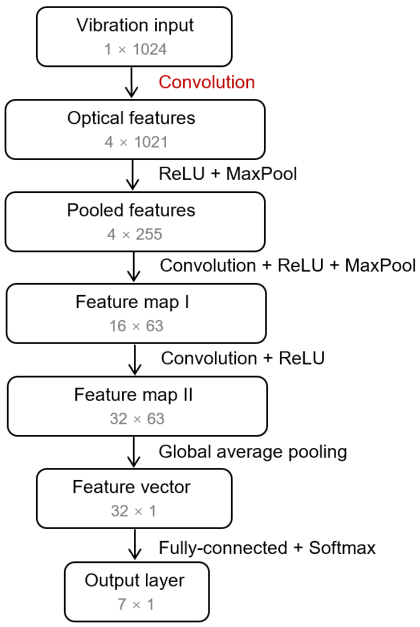

Fig. S41. Network architecture for the bearing fault diagnosis task.

The per-class confusion matrices at each precision level from 2-bit to 8-bit are shown in Fig. S42. The per-channel accuracy distribution (Fig. S43) shows that 189 of 200 optical channels achieve ≥90% accuracy, highlighting the robustness of multiplexed parallel inference across the wavelength, mode and RF dimensions.

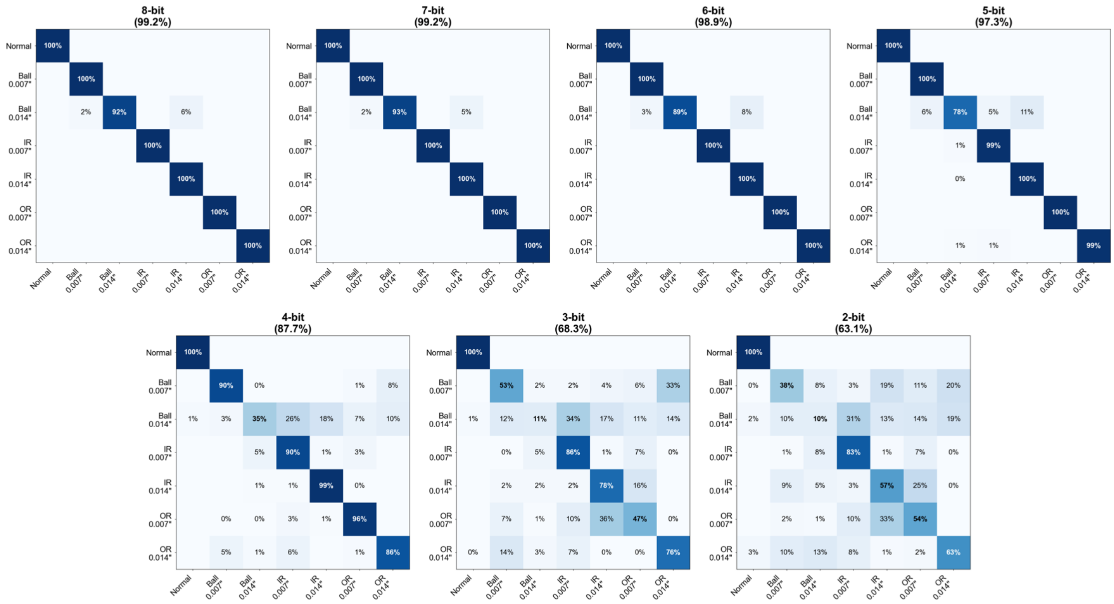

Fig. S42. Per-class confusion matrices at precision levels from 8-bit to 2-bit. Class abbreviations: IR, inner race fault; OR, outer race fault.

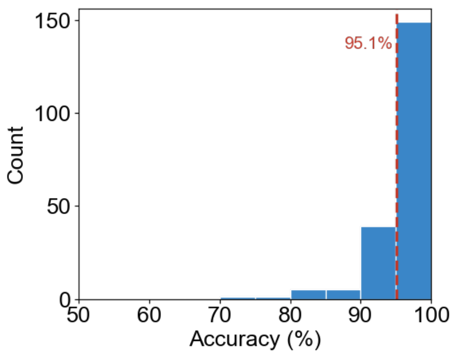

Fig. S43. Per-channel classification accuracy distribution across all 200 FieldCore optical channels. Red dashed line denotes the mean accuracy.

Supplementary Note 9: Scalability and performance estimation

A. Scaling analysis of throughput and efficiency

To evaluate the scalability and computing performance of the proposed FieldCore, we consider a fully integrated chip-level implementation in which the FieldCore is combined with the optical source, transmitter modulators, and receiver photodetectors. Each wavelength-mode pair constitutes one independent optical computing lane, and the total number of parallel optical lanes is therefore $L = N_\lambda N_m$, where N_λ is the number of wavelength channels and N_m is the number of guided-mode channels. Let R denote the input port count and K denote the output port count. Here, following the throughput definition introduced in the main-text Methods, B_{lane} denotes the loading signal rate of each wavelength-mode lane. Since one MAC corresponds to two operations, the total compute throughput can be written as

$$T_{\mathrm{comp}} = 2B_{lane}RKL = 2B_{lane}RKN_\lambda N_m. \tag{S9.1}$$

This relation shows that jointly multiplexed wavelength-mode lanes directly translate into a proportional increase in the total tensor-computing throughput. More importantly, the key architectural advantage appears at the tensor-core level. Defining the intrinsic tensor-core efficiency as

$$\eta_{\mathrm{TC}} = \frac{T_{\mathrm{comp}}}{RKP_{\mathrm{cell}}} = \frac{2B_{lane}L}{P_{\mathrm{cell}}}, \tag{S9.2}$$

where P_{cell} is the effective power consumption of one programmable computing cell, one sees that the tensor-core efficiency also increases linearly with L. In this sense, joint wavelength-guided-mode multiplexing does not merely increase the number of parallel lanes, but directly converts native optical multiplexing into power amortization of the shared programmable core.

When the optical source and electro-optic frontend power are further included, the total system power is written as

$$P_{\mathrm{tot}} = P_{\mathrm{src}} + RLP_{\mathrm{Tx}} + KLP_{\mathrm{Rx}} + RKP_{\mathrm{cell}}, \tag{S9.3}$$

where P_{src} is the source power, P_{Tx} is the total electrical power associated with one transmitter lane, P_{Rx} is the total electrical power associated with one receiver lane.

The corresponding system-level compute efficiency is

$$\eta_{\text{sys}} = \frac{T_{\text{comp}}}{P_{\text{tot}}} = \frac{2B_{lane}}{\frac{P_{\text{src}}}{RKN_\lambda N_m} + \frac{P_{\text{Tx}}}{K} + \frac{P_{\text{Rx}}}{R} + \frac{P_{\text{cell}}}{N_\lambda N_m}}. \tag{S9.4}$$

This further shows that increasing $N_\lambda N_m$ progressively amortizes both the source power and the core power, thereby continuously enhancing the system-level efficiency and driving it toward a favorable asymptotic value determined only by the transmitter and receiver frontend power per effective computing branch. The above analysis assumes a FieldCore architecture enabled by passive inverse-designed photonic components, which allow the multiplexing and demultiplexing operations to be carried out with essentially no continuous tuning or biasing power overhead.

B. Performance estimation and SOTA comparison

We next estimate the compute performance of the present FieldCore using experimentally grounded parameters. The current implementation is based on a 4×4 programmable crossbar, corresponding to $R = 4$ input ports and $K = 4$ output ports. Along the wavelength dimension, up to 36 optical carriers can be accommodated within the 1530-1565 nm window of the C band using a 1-nm spacing for interference-free optical accumulation, consistent with the experimentally validated arithmetic fidelity at representative wavelengths across this span. Since each 4-element input vector occupies four wavelengths, this corresponds to $N_\lambda = 9$ effective WDM channels. The following estimates refer to two operating regimes: a parallelism-maximized regime and a throughput-maximized regime.

In the parallelism-maximized regime, combined with dual-mode operation ($N_m = 2$) and $N_f = 100$ RF subcarriers supported in each wavelength-mode lane, the maximum number of parallel data streams is

$$N_\parallel = N_\lambda N_m N_f = 9 \times 2 \times 100 = 1800. \tag{S9.5}$$

This result shows that the RF dimension enhances parallelism by partitioning each wavelength-mode lane into reconfigurable parallel data streams.

In the throughput-maximized regime, we consider WDM-MDM operation at the full per-lane symbol rate, utilizing the entire available bandwidth without RF subcarrier partitioning. Having experimentally verified the feasibility of computation with ultrahigh-speed 120-GBaud input signals, we estimate the aggregate compute throughput as

$$T_{\text{comp}} = 2B_{\text{lane}}RKN_{\lambda}N_m = 2 \times 120 \times 4 \times 4 \times 9 \times 2 = 69.12 \text{ TOPS}. \quad \text{(S9.6)}$$

This estimate highlights the high-throughput capability of fully multiplexed photonic computing. Using the present chip footprint of $2.55 \times 2.58\ \text{mm}^2$, corresponding to a total area of $6.579\ \text{mm}^2$, the corresponding compute density is estimated to be

$$D_{\text{comp}} = \frac{T_{\text{comp}}}{A_{\text{chip}}} = \frac{69.12}{6.579} = 10.51\ \text{TOPS/mm}^2. \quad \text{(S9.7)}$$

We further estimate the system efficiency by incorporating the power contributions of the optical source, the transmitter driver-modulator frontend, the photodetector-TIA receiver frontend, and the tuning power of the photonic computing core. First, the tuning power of the photonic core is estimated from the measured thermal tuning efficiency of the MZI ($P_\pi = 28.5$ mW). For a thermo-optic MZI[16] programmed to a power coupling ratio κ, the corresponding static heating power is written as $P(\kappa) = \frac{2P_\pi}{\pi}\arcsin\sqrt{\kappa}$. Accordingly, considering the 3-dB biasing power, the static power of each multiplier is given by $P_{\text{mul}} = P\left(\frac{1}{2}\right) = \frac{P_\pi}{2} = 14.25$ mW, and the 16 multipliers therefore contribute $228\ \text{mW}$ in total. Taking the coupling ratios into the standard thermo-optic MZI relation, the total power consumptions of coupler and adder arrays are both estimated to be 139.67 mW. The total core power is thus

$$P_{\text{core}} = 228 + 139.67 + 139.67 = 507.34\ \text{mW}. \quad \text{(S9.8)}$$

For the optical source, we use a 1 W estimate based on an electrically pumped integrated soliton microcomb[17]. For the transmitter front end, we adopt a published 112-GBaud integrated CMOS-silicon-photonics transmitter, corresponding to a power consumption of 0.19 W per lane[18]. For the receiver front end, we use a reported 100-GBaud optical receiver based on a SiGe BiCMOS traveling-wave EIC and a silicon photonic Ge photodetector, with a reported power consumption of 0.5 W per lane[19]. Since the numbers of transmitter and receiver lanes are both $RL = KL = 4 \times 18 = 72$, the total system power is estimated as

$$P_{\rm sys} = P_{\rm src} + 72P_{\rm Tx,lane} + 72P_{\rm Rx,lane} + P_{\rm core}. = 51.19\ W \qquad \text{(S9.9)}$$

Accordingly, the system efficiency of fully integrated FieldCore is estimated to be

$$\eta_{\rm sys} = \frac{T_{\rm comp}}{P_{\rm sys}} = \frac{69.12\ TOPS}{51.19\ W} = 1.35\ \text{TOPS/W}. \qquad \text{(S9.10)}$$

And the photonic-core efficiency, defined by considering only the power consumption of the crossbar array, is estimated as

$$\eta_{\rm core} = \frac{T_{\rm comp}}{P_{\rm core}} = \frac{69.12\ TOPS}{0.507\ W} = 136.24\ \text{TOPS/W}. \qquad \text{(S9.11)}$$

Finally, as summarized in Table S1, state-of-the-art integrated photonic processors have so far exploited only subsets of the available degrees of freedom of the optical field, most commonly in two- or three-dimensional combinations. By contrast, the present FieldCore fully harnesses space, time, RF, wavelength and guided mode within the same photonic computing framework. Under this comparison, the present work reaches the highest dimensionality among the integrated photonic computing systems, while also combining high parallel-channel count, high operating data rate and high aggregate compute capability. Here, the parallelism is defined as the number of simultaneously processed multiplexed channels, given by the product of RF, wavelength and mode channels.

Table S1. Comparison of multidimensional integrated photonic computing architectures.

Work/Year	Dimensionality	Parallelism	Compute Speed (TOPS)	Input Data Rate (GBaud)	Compute Density (TOPS/mm²)	Energy Efficiency (TOPS/W)
Xu et al., 2022[20]	3 (T, λ, S)	4	0.48	20	0.588	/
Bai et al., 2023[21]	2 (T, λ)	4	0.136	17	1.04	2.38 [a]
Yin et al., 2023[22]	3 (T, λ, M)	12	0.537	22.38	1.37	223.2 [b]
Dong et al., 2023[7]	3 (λ, S, RF)	100	/	0.0026 [c]	0.37	0.54 [a]
Hu et al., 2025[23]	2 (T, S)	1	1.36	43.8	/	17.36 [a]
Khaled et al., 2025[24]	3 (T, λ, M)	4	/	5 [d]	/	/
Ou et al., 2025[25]	3 (T, λ, S)	7	0.98	10	0.0175	6.25 [a]
Meng et al., 2025[26]	3 (T, λ, RF)	12	0.245	30.67	34.04	0.267 [a]
Wang et al., 2026[27]	3 (T, λ, S)	5	4	50	0.16	12.5 [b]
Wu et al., 2026[28]	3 (T, λ, S)	4	0.96	30	0.0245	0.3 [a]
This work	**5 (T, λ, M, RF, S)**	**1,800**	**69.12**	**120**	**10.51**	**1.35 [a]** **136.24 [b]**

T, time; λ, wavelength; M, mode; RF, radio frequency; S, space.

[a] System level estimate; [b] estimate considering only the photonic computing core; [c] reported as the RF input frequency at 2.6 MHz; not a baudrate; [d] reported as the input non-return-to-zero (NRZ) bit rate.

Supplementary Note 10: Equipment models

The main equipment used in this paper along with their model types are listed in Supplementary Table S2.

Table S2 List of instruments and components used in the measurement and characterization setups

Equipment	**Model**
Arbitrary Waveform Generator	Keysight M8199B Keysight M8194A
Oscilloscope	Keysight UXR0594BP Agilent DSA-X 92004A
Optical Spectrum Analyzer	Yokogawa AQ6370D
Photodiode	Finisar XPDV4121R
Mach-Zehnder Modulator	NOEIC MZ135-LN-110 Sumicem T.MXH 1.5-20PD-ADC-LV
Tunable Laser	Keysight N7714A Keysight 8164B OVLINK TSP-1000
Amplified Spontaneous Emission Source	OVLINK ASE-CL-PM
Erbium-Doped Fiber Amplifier	OVLINK EDFA-C-BA-GF-26-PM-B OVLINK EYDFA-C-HP-BA-30-PM-B Amonics AEDFA-23-B-FA
DC Source	MCVS6400-A Keithley 2400
Optical Band-Pass Filter	EXFO XTM-50-SCL-U

References

1. Sun, A. *et al.* Edge-guided inverse design of digital metamaterial-based mode multiplexers for high-capacity multi-dimensional optical interconnect. *Nat Commun* **16**, 2372 (2025).
2. Deng, X. *et al.* Inverse Design of a Wavelength (De)Multiplexer for 1.55- and 2-μm Wavebands by Using a Hybrid Analog-Digital Method. *Journal of Lightwave Technology* **42**, 5231–5240 (2024).
3. Chang, W. *et al.* Ultracompact dual-mode waveguide crossing based on subwavelength multimode-interference couplers. *Photon. Res., PRJ* **6**, 660–665 (2018).
4. Wu, X. *et al.* Low Crosstalk Bent Multimode Waveguide for On-chip Mode-Division Multiplexing Interconnects. in *2018 Conference on Lasers and Electro-Optics (CLEO)* 1–2 (2018).
5. HEAT solver introduction. *Ansys Optics* https://optics.ansys.com/hc/en-us/articles/360034917713-HEAT-solver-introduction.
6. Siew, S. Y. *et al.* Review of Silicon Photonics Technology and Platform Development. *J. Lightwave Technol.* **39**, 4374–4389 (2021).
7. Dong, B. *et al.* Higher-dimensional processing using a photonic tensor core with continuous-time data. *Nat. Photon.* **17**, 1080–1088 (2023).
8. Feldmann, J. *et al.* Parallel convolutional processing using an integrated photonic tensor core. *Nature* **589**, 52–58 (2021).
9. Zhang, J. *et al.* Rate-Flexible Single-Wavelength TFDM 100G Coherent PON Based on Digital Subcarrier Multiplexing Technology. in *2020 Optical Fiber Communications Conference and Exhibition (OFC)* 1–3 (2020).
10. Xing, S. *et al.* Seamless optical cloud computing across edge-metro network for generative AI. *Nat Commun* **16**, 6097 (2025).
11. Wang, Z., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**, 600–612 (2004).
12. Luo, L.-W. *et al.* WDM-compatible mode-division multiplexing on a silicon chip. *Nat Commun* **5**, 3069 (2014).
13. Baumgardner, M. F., Biehl, L. L. & Landgrebe, D. A. 220 Band AVIRIS Hyperspectral Image Data Set: June 12, 1992 Indian Pine Test Site 3. https://doi.org/10.4231/R7RX991C doi:10.4231/R7RX991C.
14. Mourgias-Alexandris, G. *et al.* Noise-resilient and high-speed deep learning with coherent silicon photonics. *Nat Commun* **13**, 5572 (2022).
15. Bearing Data Center | Case School of Engineering. https://engineering.case.edu/bearingdatacenter.
16. Ribeiro, A. *et al.* Column-Row Addressing of Thermo-Optic Phase Shifters for Controlling Large Silicon Photonic Circuits. *IEEE J. Select. Topics Quantum Electron.* **26**, 1–8 (2020).
17. Raja, A. S. *et al.* Electrically pumped photonic integrated soliton microcomb. *Nat Commun* **10**, 680 (2019).
18. Li, K. *et al.* An integrated CMOS–silicon photonics transmitter with a 112 gigabaud transmission and picojoule per bit energy efficiency. *Nat Electron* **6**, 910–921 (2023).
19. Declercq, J. *et al.* A 100 GBd PAM-4 optical receiver using a SiGe BiCMOS traveling-wave EIC and a silicon photonic Ge photodetector. *IET Conf. Proc.* **2023**, 1378–1381 (2023).

20. Xu, S., Wang, J., Yi, S. & Zou, W. High-order tensor flow processing using integrated photonic circuits. *Nat Commun* **13**, 7970 (2022).
21. Bai, B. *et al.* Microcomb-based integrated photonic processing unit. *Nat Commun* **14**, 66 (2023).
22. Yin, R. *et al.* Integrated WDM-compatible optical mode division multiplexing neural network accelerator. *Optica, OPTICA* **10**, 1709–1718 (2023).
23. Hu, Y. *et al.* Integrated lithium niobate photonic computing circuit based on efficient and high-speed electro-optic conversion. *Nat Commun* **16**, 8178 (2025).
24. Khaled, A., Aadhi, A., Huang, C., Tait, A. N. & Shastri, B. J. Fully integrated hybrid multimode-multiwavelength photonic processor with picosecond latency. *Nat Commun* **17**, 28 (2025).
25. Ou, S. *et al.* Hypermultiplexed integrated photonics–based optical tensor processor. *Science Advances* **11**, eadu0228 (2025).
26. Meng, X. *et al.* High-integrated photonic tensor core utilizing high-dimensional lightwave and microwave multidomain multiplexing. *Light Sci Appl* **14**, 27 (2025).
27. Wang, J. *et al.* Microcomb-enabled parallel self- calibration optical convolution streaming processor. *Light Sci Appl* **15**, 149 (2026).
28. Wu, Y. *et al.* Integrated photonic 3D tensor processing engine. *Light Sci Appl* **15**, 154 (2026).