

Supplementary Information:
Broadband Control of Light through Complex Media via Automatically Self-Referencing Transmission Matrix Characterisation

M. Maestre Morote^{*1}, M. Mounaix¹, A. Komonen¹, M. Plöschner¹, and J. Carpenter¹

¹*School of Electrical Engineering and Computer Science, The University of Queensland, Brisbane, QLD 4072, Australia.*

Supplementary video: YouTube link

1. Motivation - 0s
2. Multimode fibre characterisation - 1m:33s
3. Stokes polarimetry analogy - 2m:48s
4. Spatial state tomography - 6m:0s
5. Mode transmission matrix retrieval - 8:25s
6. Long coherence source results - 9m:50s
7. Short coherence source results - 12m:21s
8. Conclusions - 15m:57s

Supplementary code:  [GitHub](#) [MarKo7s/SST MMF characterization](#) [release](#) [v0.1.0 arXiv snapshot](#) [DOI](#) [10.5281/zenodo.21637820](#)

Processing routines and data to retrieve the MTM of a simulated GI fibre and of the experimental OM1 fibre (five mode groups, five $\Delta\lambda_S$), Version 0.1.0 (arXiv snapshot); DOI 10.5281/zenodo.21637820. Notebooks: simulation, experimental.

Scope

Evaluation metrics (overlap integral OI , singular-value analysis, pattern generation, and SCS emulation) and phase locking between camera pixels are defined in the main article Methods. This Supplementary Information provides:

- Note 1 — SST theory;
- Note 2 — MMF characterisation and MTM assembly;
- Note 3 — extended laboratory setup, alignment, and data handling;
- Note 4 — additional LCS figures with interpretation;
- Note 5 — fibre correlation bandwidth Λ ;
- Note 6 — secondary eigenstates and extended SCS beam shaping.

^{*}corresponding author: m.maestremorote@uq.edu.au

Supplementary Note 1: Spatial state tomography background

We provide a brief overview of state tomography using quantum formalism before plunging into the parallelised high-dimensional spatial tomography used to characterise multimode fibre. Spatial state tomography (SST) is a generalisation of state tomography applied to the spatial degree of freedom of light. It can be visualised as the high-dimensional counterpart to Stokes polarimetry.

In this section we index fibre modes by $n \in \{0, \dots, N-1\}$ in the basis $\{|\psi_n\rangle\}$; density-matrix eigenstates $|\varphi_i\rangle$ by $i \in \{1, \dots, N\}$; Gell–Mann matrices and Stokes parameters by $j \in \{0, \dots, N^2-1\}$ ($\hat{\sigma}_j, S_j$); the unique eigenstates of each $\hat{\sigma}_j$ by r ($|\Omega_j^r\rangle, \kappa_j^r$), when referenced to a particular j -matrix or in a flat indexing format by $k \in \{1, \dots, m\}$ with $m = 2N^2 - N$.

Density-matrix formalism

An unknown state of light, which could be a pure state or a mixed state (i.e., an incoherent superposition of multiple pure states), is described by a density matrix $\hat{\rho}$. By the spectral theorem, it can be represented as its eigendecomposition:

$$\hat{\rho} = \sum_{i=1}^N p_i |\varphi_i\rangle \langle \varphi_i|, \quad (\text{S1})$$

where p_i is the probability (or intensity weighting) of finding a given $|\varphi_i\rangle$. The pure eigenstates $|\varphi\rangle$ in the mixture $\hat{\rho}$ are linear combinations of the N modes forming the basis $\{|\psi_n\rangle\}$ along the degree of freedom in which the tomography is performed:

$$|\varphi\rangle = \sum_{n=0}^{N-1} c_n |\psi_n\rangle, \quad (\text{S2})$$

where c_n are complex coefficients.

To reconstruct $\hat{\rho}$, a tomographically complete set of observables must be measured. For an N -level system these are the generalised Gell–Mann matrices $\hat{\sigma}_j$: $N^2 - 1$ linearly independent $N \times N$ traceless Hermitian matrices that, together with the identity, span the space of density matrices. With normalisation $\text{Tr}(\hat{\sigma}_j \hat{\sigma}_\ell) = 2\delta_{j\ell}$, the density matrix is recovered as a weighted sum:

$$\hat{\rho} = \frac{1}{2} \sum_{j=0}^{N^2-1} S_j \hat{\sigma}_j, \quad (\text{S3})$$

where $S_j = \text{Tr}(\hat{\rho} \hat{\sigma}_j)$ are the Stokes parameters collected in the high-dimensional Stokes vector $\mathbf{S}$. How these expectation values are obtained experimentally is developed below.

Observable expectation, Gell–Mann matrices, projection operator and Stokes vector

Projective measurements are the specific outcomes used to infer physical properties of an unknown system; in the tomographic formalism these quantities are called observables. An observable is a linear Hermitian operator $\hat{O}$ acting on the N -dimensional Hilbert space. For a mixed state $\hat{\rho}$, the expectation value follows the Born rule:

$$\langle O \rangle_{\hat{\rho}} = \sum_{i=1}^N p_i \langle \varphi_i | \hat{O} | \varphi_i \rangle = \text{Tr}(\hat{\rho} \hat{O}), \quad (\text{S4})$$

where Tr denotes the trace. The Gell–Mann matrices $\hat{\sigma}_j$ generalise the Pauli matrices from $\text{SU}(2)$ to $\text{SU}(N)$, forming an orthogonal basis for the Lie algebra of traceless Hermitian operators.

The expectation values S_j of finding an observable $\hat{\sigma}$ in the mixture is not always directly measurable. Instead each $\hat{\sigma}_j$ is spectrally decomposed into projection operators onto its eigenstates, the analyser states $|\Omega_j^r\rangle$:

$$\hat{\sigma}_j = \sum_r \kappa_j^r |\Omega_j^r\rangle \langle \Omega_j^r| = \sum_r \kappa_j^r \hat{P}_j^r, \quad (\text{S5})$$

where κ_j^r are the corresponding eigenvalues and $\hat{P}_j^r = |\Omega_j^r\rangle\langle\Omega_j^r|$ the projection operator for the eigenvectors. Projecting each analyser state onto $\hat{\rho}$ yields the Stokes component:

$$S_j = \text{Tr}(\hat{\rho} \hat{\sigma}_j) = \sum_r \kappa_j^r \langle\Omega_j^r| \hat{\rho} |\Omega_j^r\rangle. \quad (\text{S6})$$

Similarly, if an observable comprises several eigenstates, each must be projected individually to remove ambiguity. One of each of the Gell–Mann matrices therefore explores a subspace of the $N \times N$ Hilbert space and projecting its eigenstates one-by-one provides unambiguous intensity outcomes. The detected intensity $I_{|\Omega_k\rangle} = \langle\Omega_k| \hat{\rho} |\Omega_k\rangle$ per analyser state is equivalent to the integral overlap over the detection area (A) between all the fields inside the mixture ($\varphi_i(A)$) and the projected field ($\Omega_k(A)$):

$$I_{|\Omega_k\rangle} = \langle\Omega_k| \hat{\rho} |\Omega_k\rangle = \sum_i p_i \langle\Omega_k| \varphi_i \rangle \langle\varphi_i| \Omega_k\rangle = \sum_i p_i |\langle\Omega_k| \varphi_i \rangle|^2 = \sum_i p_i \left| \int \Omega_k^*(A) \varphi_i(A) dA \right|^2 \quad (\text{S7})$$

Finally another way to visualise the Stokes vector $\mathbf{S}$ is as the Cartesian coordinates of $\hat{\rho}$ in the Gell–Mann basis (Eq. (S3)), living in a real hypersphere of dimension $N^2 - 1$. Purity sets the radius: the maximally mixed state is at the centre ($\mathbf{S} = 0$), pure states lie on the surface, and partially mixed states in between. For $N = 2$ this is exactly the Poincaré sphere, with (S_1, S_2, S_3) as the familiar polarisation Stokes parameters (following subsection).

Tomographic setup and efficient reconstruction

The physical measurement requires projecting a total of $m = 2N^2 - N$ unique analyser states $|\Omega_k\rangle$ onto the unknown state. These are the eigenvectors of $\hat{\sigma}_j$, selected with a uniqueness threshold $\Theta = 1 - 1/N^2$: a new eigenstate $|\Omega_j^r\rangle$ is retained as $|\Omega_k\rangle$ only if it is sufficiently distinct from all previously stored states ($\langle\Omega_k| \Omega_j^r\rangle < \Theta$). The associated eigenvalues κ_j^r populate a sparse weighting matrix $\mathbf{K}$ of dimension $N^2 \times (2N^2 - N)$, with $\mathbf{K}_{jk} = \kappa_j^r$ when analyser state k corresponds to the r -th eigenstate of $\hat{\sigma}_j$, and zero otherwise.

After acquiring the intensity of each projection, the Stokes vector is recovered by matrix multiplication:

$$\mathbf{S} = \mathbf{K} \cdot \mathbf{I}_\Omega, \quad (\text{S8})$$

where $\mathbf{I}_\Omega(x, y)$ is the intensity vector with entries $I_{|\Omega_k\rangle} = \langle\Omega_k| \hat{\rho} |\Omega_k\rangle$. Formulating the problem with help of Eq. (S8) extends $\mathbf{S}$ to parallel SST problems by stacking independent $\mathbf{I}_\Omega$ vectors columns and performing the matrix multiplication.

Non-ideal tomography

The framework above assumes ideal conditions: a perfectly matched basis, projections executed with 100% accuracy, noiseless detection, and a static system throughout the acquisition. In practice, imperfections introduce errors that fall into three categories: (1) errors in the measurement basis, (2) counting and statistical errors, and (3) experimental-stability errors such as source-intensity, mechanical, or thermal drift. Any of these perturb the measured expectation values, so the retrieved state always deviates somewhat from the true one; in the Stokes-polarimetry picture the solution is no longer a single point on the Poincaré sphere but lies within an error uncertainty “cloud of points”.

A rigorous error analysis is non-trivial. Conventionally it relies on repeating the tomography many times and fitting each error source to its statistical distribution via maximum-likelihood estimation (MLE) or Bayesian methods, which both require a model per error source and scales poorly, as the number of projections grows rapidly with the system dimension. Even then, the recovered matrix is not guaranteed to be physical, so a final projection onto the nearest positive semi-definite (PSD) matrix is required [1]. For our system this counting-statistics approach is impractical: the dimensionality is simply too large to repeat the measurement.

Instead, we shrink the “error uncertainty cloud” proactively through calibration—aberration correction, mask efficiency optimisation, and frame averaging where possible—and then compute the nearest PSD density matrix and threshold sub-noise eigenvalues before extracting the eigenstates. Some errors, notably mechanical and thermal drift over extended acquisitions, remain beyond calibration: when the system drifts faster than it can be measured, previously inactive states gain probability weight and the mixture artificially grows, an effect that becomes more pronounced at higher dimension. Crucially, this noise predominantly corrupts the eigenvalues, which are largely irrelevant to our application, while the

eigenvectors (the spatial patterns) remain robust. The measurement can be viewed as a power-maximisation problem: after probing with all projections, the primary eigenstate emerges as the linear superposition of basis modes that maximises transmitted power to that output channel. Drift and noise redistribute probability into additional, weakly weighted eigenstates, so the weight (p_1) of the physical solution may decrease; the modal superposition itself, however, remains essentially unchanged. That eigenvalue is only a predicted power fraction for that superposition. Because we seek the optimal focusing solution rather than an accurate absolute power forecast, this reduction in p_1 does not prevent reliable recovery of the principal eigenstate.

For these reasons we avoid measurement-reduction strategies—such as projecting only N^2 states and inferring the remainder, or compressed sensing with fewer than N^2 projections—which could shorten acquisition but introduce additional errors from power fluctuations, projection imperfections, or non-convex optimisation. We therefore project all $m = 2N^2 - N$ analyser states, accepting longer acquisition in exchange for tomographic completeness.

Stokes polarimetry analogy ($N = 2$)

For $N = 2$, SST reduces to standard Stokes polarimetry: the Gell–Mann observables are the Pauli matrices,

$$\hat{\sigma}_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \hat{\sigma}_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \hat{\sigma}_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \hat{\sigma}_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (\text{S9})$$

and $\mathbf{S} = (S_0, S_1, S_2, S_3)$ gives Cartesian coordinates on the Poincaré sphere (Eq. (S3)). Each $\hat{\sigma}_j$ decomposes as in Eq. (S5); for $N = 2$ the $m = 6$ analyser states are the familiar polarisation bases ($|H\rangle, |V\rangle, |D\rangle, |A\rangle, |R\rangle, |L\rangle$), with orthogonal eigenpairs on opposite Poincaré-sphere axes. Table S1 lists eigenvalues, eigenvectors, analyser states, and the Stokes components from Eq. (S6).

Table S1: Mathematical representation of an $N = 2$ state tomography setup.

Pauli matrix from Eq. (S9)	eigenvalues as $\hat{\sigma}_j = U \cdot Y \cdot U^{-1}$		eigenvectors		Analyser state applying Eq. (S2)		Stokes component applying Eq. (S6)
	κ_j^0	κ_j^1	v_j^0	v_j^1	$ \Omega_j^0\rangle$	$ \Omega_j^1\rangle$	S_j
$\hat{\sigma}_0$	1	1	$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 0 \\ 1 \end{pmatrix}$	$ H\rangle$	$ V\rangle$	$\langle H \hat{\rho} H\rangle + \langle V \hat{\rho} V\rangle$
$\hat{\sigma}_1$	1	-1	$\begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}$	$\begin{pmatrix} \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \end{pmatrix}$	$ D\rangle$	$ A\rangle$	$\langle D \hat{\rho} D\rangle - \langle A \hat{\rho} A\rangle$
$\hat{\sigma}_2$	1	-1	$\begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{i}{\sqrt{2}} \end{pmatrix}$	$\begin{pmatrix} \frac{1}{\sqrt{2}} \\ -\frac{i}{\sqrt{2}} \end{pmatrix}$	$ R\rangle$	$ L\rangle$	$\langle R \hat{\rho} R\rangle - \langle L \hat{\rho} L\rangle$
$\hat{\sigma}_3$	1	-1	$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 0 \\ 1 \end{pmatrix}$	$ H\rangle$	$ V\rangle$	$\langle H \hat{\rho} H\rangle - \langle V \hat{\rho} V\rangle$

While the previous table sets out the theory for a generic $N = 2$ tomographic measurement, Figure S1 illustrates its experimental realisation, comparing Stokes polarimetry with the spatial $N = 2$ analogue side by side. The tomographic framework is unchanged: only the degree of freedom under measurement differs (polarisation versus space). In both, (a) and (b), the same unknown input mixed state, denoted $\hat{\rho}_{in}$, is prepared in the basis appropriate to each degree of freedom. The six analyser states are projected sequentially using a $\lambda/2$ plate, a $\lambda/4$ plate, and a polariser in conventional Stokes polarimetry, and using an SLM in SST. For each analyser state, the transmitted intensity is recorded with a single-pixel detector formed by a lens, single-mode fibre (SMF), and photodiode, equivalently to Eq. (S7). Once all projections are complete, the mixed state is reconstructed following the formalism above: the measured intensities are converted into the Stokes vector $\mathbf{S}$ (c and e), $\hat{\rho}_{in}$ is recovered from $\mathbf{S}$ (d), and its eigenstates are extracted by eigendecomposition (f).

Finally, (g) and (h) show that the unknown mixture is an incoherent superposition of two pure states with intensity weights 3:2—the same mixture encoded in polarisation and in space.




Figure S1: Stokes polarimetry analogy for $N = 2$. (a, b) Experimental setups for polarisation and spatial tomography, using waveplates/polariser or an SLM respectively. (c, e) Intensity measurements are converted into a Stokes vector $\mathbf{S}$. (d, f–h) The density matrix is reconstructed and eigendecomposed to recover the pure states in the mixture. Adapted from [2].

Supplementary Note 2: SST concept for MMF characterisation

The SST concept applied to MMF characterisation remains the same but extends the dimensionality to the number of spatial modes supported by the fibre-under-test. For the particular case presented in the main article, the employed fibre is an OM1 GI-MMF supporting 20 mode groups at 1300 nm and two polarisation states. This is equivalent to 420 modes across both polarisations (210 each) since they are coupled by the fibre. Analogously to the previous SST example for $N = 2$, this scenario is expanded as shown in Figure S2.




Figure S2: Input analyser states for MMF characterisation. (a) Six analyser states for Stokes polarimetry ($N = 2$). (b) High-dimensional SST for the OM1 GI-MMF ($N = 420$), requiring $2N^2 - N = 352,380$ pairwise mode-interference projections at four phase offsets, plus single-mode terms.

A distinctive feature of MMF characterisation employing SST is that the fibre itself does not constitute a mixed state, it produces a wavelength dependent spatial transformation to the input light. For a coherent source at a single wavelength, the output is a pure state—a linear combination of the supported basis modes weighted by fibre-induced mode mixing, which varies with input mode and modal dispersion. It can be inferred that if we can probe the fibre with a set of orthogonal basis at the input and untangle the mixing using the above SST formalism we can construct a TM row by row.

A natural probing basis is the spot basis: M^2 illuminations on the proximal facet, arranged for simplicity as an $M \times M$ array (Figure S3(a)). For each input spot (x, y) , the fibre produces an unknown output field at the distal facet. Applying the same SST protocol as in the Stokes analogy—analyser-state projection with an SLM and readout by a single-pixel detector (lens, SMF, photodiode)—recovers the modal superposition at the output, i.e. the complex coefficients that map that input spot to the distal field (or, for a partially coherent source, the incoherent mixture of modal superpositions). Repeating for all M^2 spots yields a TM mapping N input modes to M^2 output spots, albeit phase-unresolved between spots. The cost is prohibitive: each row requires a full $2N^2 - N$ projection sequence, so the conventional sequential approach scales as $M^2 \times (2N^2 - N)$ acquisitions.

We invert the problem by exploiting optical reciprocity and parallel detection with a single camera (Figure S3(b)). Instead of analysing the output field separately for each input spot, the SLM injects the $2N^2 - N$ analyser states at the proximal facet; one camera frame per analyser state records the intensity at every pixel $I(x, y, k)$. The pixel grid provides an orthogonal spatial sampling basis: each pixel measures an intensity expectation in a distinct region of the output plane, so M^2 virtual tomographies run in parallel. Processing the k -stack at each pixel (x, y) yields a local density matrix $\hat{\rho}_{x,y}$ and its primary eigenstate—the input modal superposition that focuses power at that pixel, as described in the main article. The same $N \leftrightarrow M^2$ mapping as in Figure S3(a) is obtained in a single tomographic pass rather than M^2 sequential measurements.

MTM calculation

Figure S4 highlights a key feature of the retrieval pipeline: each camera frame is split into horizontal and vertical components by the distal Wollaston prism, and the two halves are stored and processed independently. This polarisation diversity




Figure S3: SST as an MMF characterisation technique. (a) Each row of the TM is obtained by performing high-dimensional tomography at the fibre output for M^2 orthogonal input illuminations. (b) Injecting the $2N^2 - N$ analyser states at the input and recording $I(x, y, k)$ with a camera parallelises M^2 virtual tomographies, yielding one density matrix per pixel.

yields two spot-domain transmission matrices, TM_H and TM_V , each mapping the N input modes— $N/2$ spatial modes per polarisation—to the M^2 camera pixels.




Figure S4: Polarisation-resolved data layout for MTM retrieval. Each acquired frame is divided into horizontal and vertical regions by the Wollaston prism. The two regions of interest (ROIs) are processed independently to build TM_H and TM_V in the output spot basis.

Each spot-domain matrix is initially known only up to an unknown phase per pixel. The phase-locking algorithm described in the main manuscript (Methods A) recovers the relative phase between neighbouring pixels within the same speckle grain, propagating a globally consistent phase across the full M^2 spot basis without an external reference. Once phase-corrected, each TM is transformed from the output pixel basis to the Laguerre–Gaussian (LG) input basis of dimension $N/2$, giving MTM_H and MTM_V . Interleaving the two polarisation blocks assembles the complete $N \times N$ polarisation-resolved MTM that describes coupled spatial and polarisation propagation through the fibre (Figure S5). This eliminates the redundant oversampling of the spot representation and expresses the measured response entirely in the modal basis used for wavefront shaping. As detailed in [3, 4], the condition $M^2 \geq 4N$ is sufficient to retrieve the complete MTM; downsampling of the camera grid can be applied when this oversampling criterion is met.




Figure S5: Complete polarisation-resolved MTM assembly. For each polarisation (H, V), the phase-corrected spot-domain TM is transformed to the LG input basis. The two blocks are interleaved into a single $N \times N$ MTM that fully describes spatial and polarisation coupling through the MMF.

Supplementary Note 3: Laboratory diagram, alignment and data handling

Figure S6 expands Fig. 1(a) of the main text with colour-coded stages. The apparatus is designed for two characterisation stages on the same optical bench: self-referencing SST, in which the holographic reference arm is blocked and only distal intensities are recorded; and off-axis digital holography (DH), in which the reference arm is unblocked to recover complex output fields for benchmarking. Evaluation metrics, the swept-laser protocol, and phase locking are defined in the main article Methods B, D, and A. This note records the alignment practice, acquisition layout, and high- N storage and processing workflow required to reproduce the measurements.




Figure S6: Extended laboratory setup. Blue: proximal wavefront-shaping stage (SLM, polarisation optics, and lenses). Pink: distal detection with Wollaston prism for polarisation diversity. Green: swept-laser trigger path. Left beam blocker (BB): off-axis digital holography reference arm; right BB: SST intensity-only detection.

Optical stages

The proximal stage (blue) generates arbitrary spatial and polarisation fields at the fibre input. A swept-frequency laser (New Focus Venturi TLB-8800, 80 nm tuning range centred at 1300 nm) feeds an LCoS SLM (Meadowlark P1920) placed near the Fourier plane of the proximal optics. Waveplates, polarisation beam splitters, and a Gerchberg–Saxton algorithm provide independent horizontal (H) and vertical (V) control of the launched field. The fibre-under-test (displayed in orange) is a 2 m OM1 graded-index MMF (62.5 μm core) supporting 20 near-degenerate mode groups at 1300 nm, equivalent to $N = 420$ spatial and polarisation modes. The distal stage (pink) images the fibre facet onto a Xenics Cheetah InGaAs camera (640 $\times$ 512 pixels, 20 μm pitch). A calcite Wollaston prism (CWP) separates H and V onto opposite halves of the sensor so that each polarisation is processed independently. The green path implements swept-laser interferometry: a balanced photodiode converts the interferogram fringes into TTL triggers that synchronise wavelength-resolved camera frames during SCS emulation (main Methods, swept-laser protocol) [5].

Proximal alignment

Accurate proximal coupling is critical because the analyser states are constructed in the LG modal basis; misalignment or uncorrected aberrations degrade mask fidelity and propagate into every projection. Alignment proceeds in two steps with the MMF temporarily replaced by a single-mode fibre (SMF) and a power meter at the proximal facet.

First, a blazed grating displayed on the SLM deflects the zero-order away from the fibre and couples only the first diffraction order. The SMF position is adjusted manually until maximum throughput is achieved, placing the fibre as close as practicable to the SLM Fourier plane while minimising residual beam tilt that cannot be corrected digitally. Higher diffraction orders are rejected by the fibre cladding (single lens system) or, where used, a pinhole ($4f$ configuration system).

Second, residual aberrations from the SLM optics and any manual misalignment are corrected by optimising Zernike polynomials (degree ≤ 6) displayed on the SLM while monitoring coupled power into the SMF. Optimisation runs until convergence; the MMF is then reinstalled without moving the SLM carriage or proximal optics, preserving the calibrated wavefront.

Distal end and dual detection paths

For SST measurements the reference arm is blocked and only intensity is recorded at the camera. No distal phase correction is required because the tomographic reconstruction and phase-locking algorithm (main Methods A) recover the complex TM from intensity data alone. For DH benchmarking the reference arm is unblocked, forming off-axis interferograms on both polarisation halves of the sensor. Residual aberrations in the reference and imaging path are removed in software following the calibration procedures in [6, 7]. Switching between SST and DH therefore requires only repositioning the beam block; the fibre, camera, and polarisation optics remain fixed.

Tomographic data layout

Raw SST data are organised as a three-dimensional intensity array $I(x, y, k)$, where k indexes the $m = 2N^2 - N$ input analyser states and (x, y) spans the camera ROI. For the full 20 mode-group dual polarisation measurement ($N = 420$), each k -frame yields two 200×200 pixel, 16-bit ROIs (one per polarisation), giving ~ 110 GB on disk for a single LCS acquisition. Frames are stored so that each pixel can be loaded independently along k , keeping RAM usage bounded during post-processing.

For SCS emulation the swept-laser trigger records $D = 118$ wavelength-resolved frames per analyser state ($\lambda = 1280\text{--}1320$ nm), equivalent to D independent monochromatic measurements acquired in one experimental run. Because the camera is much faster than the LCoS refresh rate, the total SST acquisition time is mostly set by hologram generation, not by the number of triggered wavelengths. The spectrally resolved archives scale to ~ 10 TB and are kept on network-attached storage (NAS). Emulating a desired source bandwidth $\Delta\lambda_S$ in post-processing requires reading and summing intensity frames along the wavelength axis for each k (main Methods D); in our configuration, NAS read bandwidth dominates this step (~ 19 h), whereas the on-line tomographic acquisition duration is unchanged.

Sparse reconstruction

At $N = 420$, the dense Gell–Mann operators $\hat{\sigma}_j$ and weighting matrix $\mathbf{K}$ exceed available RAM if assembled naively. Both structures are highly sparse: most $\hat{\sigma}_j$ entries are zero, and $\mathbf{K}$ contains only the eigenvalues κ_j^r associated with each retained analyser state. Exploiting this sparsity is essential for practical reconstruction of Eqs. (S3) and (S8). Sparse implementations of the Stokes-vector and density-matrix recovery are provided in the GitHub repository https://github.com/MarKo7s/SST_MMF_characterization; post-processing of the full 110 GB LCS dataset is CPU-GPU accelerated (Intel i9-9980XE and Nvidia Titan V) and completes in ~ 30 min.

Supplementary Note 4: Long coherence source — additional figures

Short acquisition-time SST

Main text Fig. 2 shows phase ripples on some LCS targets despite favourable singular values (main Methods C). We attribute this to 54 h drift, not SST bias. The results shown in Fig. S7 correspond to SST MTM measurements on 5 and 9 mode groups only (17 min for 30 modes; 2 h for 90 modes) which removes the phase artefacts seen on the long tomographic measurement.




Figure S7: Effect of acquisition time on LCS beam shaping. Partial-mode-group SST (5 and 9 mode groups; 17 min and 2 h acquisitions) yields flat reconstructed phases, whereas the full 420-mode run acquired over 54 h exhibits phase ripples attributed to mechanical and thermal drift rather than SST bias.

Intensity overlap figure

Figure S8 plots the intensity overlap OI_I alongside the complex overlap OI (defined in Methods B). OI_I tracks visual quality more closely because it neglects phase artefacts that can arise from the holographic reference arm or from mechanical and thermal drift during a long tomographic acquisition. In this comparison, MTM_{SST} outperforms MTM_{DH} for those patterns whose measured phase is clearly non-flat: the phase errors reduce the complex OI even when intensity delivery remains strong.




Figure S8: Comparison of intensity and complex overlap metrics for LCS beam shaping. (a) Intensity overlap OI_I for four target patterns, showing closer agreement between MTM_{SST} and visual field quality. (b) Target fields, theoretical best-case patterns, and DH-measured outputs; complex OI values are annotated above each measured field.

Supplementary Note 5: Fibre temporal response and correlation bandwidth

The temporal and spectral properties of the tested OM1 fibre can be examined by performing the Fast Fourier Transform (FFT) across the frequency/wavenumber axis of the spectrally resolved MTM, which was measured using an external reference and off-axis digital holography. By grouping the modes into near-degenerate mode groups and computing the FFT, the impulse response of each mode group is obtained, revealing the time delay associated with each group.

Figure S9 illustrates the spectral and temporal results for the OM1 fibre. The 2 m custom patch cord exhibits a very narrow total impulse response, meaning each individual mode group travels very close to each other with minimal inter-group modal dispersion. This temporal behaviour translates into a broad correlation bandwidth in the frequency domain, with a total fibre correlation bandwidth $\Delta \approx 10$ nm.




Figure S9: Spectral and temporal characterisation of the 2 m OM1 GI-MMF from a spectrally resolved MTM measured by off-axis digital holography. (a) Normalised impulse response of each near-degenerate mode group, showing minimal inter-group delay. (b) Generalised frequency correlation of the full transmission matrix, yielding a fibre correlation bandwidth $\Delta \approx 10$ nm used to interpret the SCS results in the main text and Note 6.

Supplementary Note 6: Short coherence source — additional figures

Eigenvalue maps

Figure S10 displays the spatial distribution of the first seven eigenstates at experimentally emulated source bandwidths of $\Delta\lambda_S = 0, 10, 20,$ and 30 nm. The thresholding and positive semi-definite (PSD) projection procedures follow the methodology outlined in Note 1. The absolute scales are indicative, but the spatial trends clearly reflect mode-dependent loss and decorrelation. As the bandwidth increases, the primary eigenstate (φ_1) tends to concentrate power in the core (lower-order modes), while the secondary eigenstates begin to capture the decorrelated power spreading towards the cladding (higher-order modes).




Figure S10: Spatial distribution of the first seven density-matrix eigenvalues for emulated source bandwidths $\Delta\lambda_S = 0, 10, 20,$ and 30 nm. Left column: mean eigenvalue maps; right columns: per-pixel distributions. As $\Delta\lambda_S$ increases, φ_1 concentrates in the core while secondary eigenstates gain weight towards the cladding, reflecting spatial decoherence across mode groups.

Secondary MTMs

The MTM per eigenstate ($\text{MTM}\varphi_i$) are reconstructed by collapsing all pixels associated to an eigenstate into the MTM as outlined in Note 2. Figure S11 illustrates the secondary eigenstates and their corresponding singular values for $i = 1, 2, 3$.

In the monochromatic case ($\Delta\lambda_S = 0$), the system is fully spatially coherent, meaning the density matrix possesses only a single non-zero eigenstate (φ_1). Consequently, any channels appearing in $i \geq 2$ are merely experimental artefacts or noise, as discussed in the main Methods C.

However, as the source bandwidth increases, spatial decoherence emerges. For instance, at $\Delta\lambda_S = 30$ nm, the highest eigenstate (φ_1) can no longer maintain a phase relationship across all modes, causing certain higher-order mode groups to drop out of $\text{MTM}(\varphi_1)$. These "missing" modes are instead captured by the secondary eigenstates. $\text{MTM}(\varphi_2)$ and $\text{MTM}(\varphi_3)$ thus carry these sideband and cladding-weighted solutions, representing mutually incoherent components of the mixed state that can still be utilised for specific wavefront shaping tasks.




Figure S11: (i) Mode transmission matrices $\text{MTM}(\varphi_i)$ for eigenstates $i = 1, 2, 3$ and emulated bandwidths $\Delta\lambda_S = 0, 10, 20$, and 30 nm. At $\Delta\lambda_S = 0$, only φ_1 is physical; $i \geq 2$ matrices are noise. At $\Delta\lambda_S = 30$ nm, higher-order mode groups absent from $\text{MTM}(\varphi_1)$ reappear in $\text{MTM}(\varphi_2)$ and $\text{MTM}(\varphi_3)$. (ii) Normalised singular values of $\text{MTM}(\varphi_i)_{\Delta\lambda_S}$ for $i = 1, 2, 3$ and $\Delta\lambda_S = 0, 10, 20$, and 30 nm. Channel attenuation in φ_1 increases with bandwidth as fewer modes remain mutually coherent; secondary eigenstates retain usable channels for decorrelated modal subsets.

Beam shaping: Random Superposition of modes $\text{MTM}(\varphi_1)$ - Additional results

To corroborate the consistency of the broadband matrix enhancement factor discussed in the main text, Figure S12 shows the performance when generating a random amplitude and phase superposition of all modes ($\delta\lambda_T \approx \Lambda$). The results follow the exact same trend as the smiley/sad pattern, confirming the $\sqrt{2}$ bandwidth enhancement factor when using $\text{MTM}_{2\Lambda}$.

Broadband beam shaping ($\Delta\lambda_S = 30$ nm) with $\text{MTM}(\varphi_{1,2,3})$

Figure S13 explores the practical utility of the secondary eigenstates by performing beam shaping (smiley/sad, random, and off-centre spots) using $\text{MTM}(\varphi_{1,2,3})$ at $\Delta\lambda_S = 30$ nm. The results reveal a clear spatial segregation of control. $\text{MTM}(\varphi_1)$ heavily favours centre-band control (lower-order modes), while $\text{MTM}(\varphi_2)$ and $\text{MTM}(\varphi_3)$ control the sidebands and higher-order modes. Notably, when attempting to focus a spot near the cladding (which requires higher-order




Figure S12: Broadband beam shaping with a random amplitude-phase superposition of all fibre modes. Overlap $OI(\lambda)$ and integrated $OI(\Delta\lambda)$ for each retrieved $MTM_{\Delta\lambda_S}$, reproducing the $\sqrt{2}$ bandwidth-enhancement trend observed for the smiley/sad target in the main text. Optimal performance occurs when the shaping matrix matches the source bandwidth.

modes with shorter correlation bandwidths), the performance of $MTM(\varphi_3)$ improves relative to its performance on centre-focused targets, confirming that these secondary eigenstates capture the specific modal information that was "lost" from the primary eigenstate due to spatial decoherence.




Figure S13: Beam shaping at $\Delta\lambda_S = 30$ nm using $\text{MTM}(\varphi_1)$, $\text{MTM}(\varphi_2)$, and $\text{MTM}(\varphi_3)$. Rows: smiley/sad face, random superposition, and off-centre diffraction-limited spots at three radial positions. φ_1 controls core-weighted patterns; φ_2 and φ_3 improve sideband and near-cladding spots that rely on higher-order, shorter- Λ modes.

References

- [1] N. J. Higham, Linear Algebra and its Applications **103**, 103 (1988).
- [2] M. Plöschner, M. M. Morote, D. S. Dahl, M. Mounaix, G. Light, A. D. Rakić, and J. Carpenter, Nature Communications **13**, 4294 (2022).
- [3] M. Plöschner, T. Tyc, and T. Čižmár, Nature Photonics **9**, 529 (2015).
- [4] R. N. Mahalati, R. Y. Gu, and J. M. Kahn, Optics Express **21**, 1656 (2013).
- [5] M. Maestre Morote, Swept-source interferometer camera trigger (pySSTri) (2026), firmware Polarity Low 1.3.1; source <https://github.com/MarKo7s/swept-source-trigger-interferometer>.
- [6] J. Carpenter, digHolo : High-speed library for off-axis digital holography and Hermite-Gaussian decomposition (2022), arXiv:2204.02348 [eess].
- [7] M. Mounaix and J. Carpenter, Nature Communications **10**, 5085 (2019), number: 1 Publisher: Nature Publishing Group.