
PROXIMAL-ONLY TRANSMISSION MATRIX RECOVERY OF AN ARBITRARILY DEFORMED GRADED-INDEX MULTIMODE FIBER

PREPRINT

Cole Reynolds
Weyl Labs
Cole@WeylLabs.com

ABSTRACT

The multimode fiber is among the thinnest imaging conduits available, carrying hundreds to thousands of spatial modes through a cross-section comparable to a human hair, but its endoscopic capabilities are currently limited by the sensitivity of the transmission matrix to the fiber’s deformed state. Proximal-only recovery of the fiber’s transmission matrix is an appealing approach for enabling general use multimode fiber endoscopy, and within the last decade, machine learning techniques have been applied to both single-ended and double-ended transmission matrix recovery tasks. We present a new approach to this interdisciplinary problem and show that neural networks can generalize to recover transmission matrices of an arbitrarily deformed graded-index multimode fiber from proximal measurements alone.

1 Introduction

A Multimode fiber (MMF) propagates information between its proximal and distal ends through the reciprocal relationship

$$\psi_{\text{dist}} = U\psi_{\text{prox}}, \quad \psi_{\text{prox}} = U^T\psi_{\text{dist}} \quad (1)$$

where U is the fiber’s transmission matrix (TM). In MMF endoscopy, the fiber’s distal tip is guided to a scene described by an unknown scattering matrix S , and probe signals from the proximal end propagate through the fiber, interact with the scene, and a portion of that initial signal energy propagates back through the fiber to the proximal end which can then be measured:

$$\psi_{\text{measured}} = U^T S U \psi_{\text{probe}}. \quad (2)$$

For a known $N \times N$ TM, the available scene content S can be recovered, in principle, with N linearly independent probe signals. This is the basis of MMF endoscopy, and it is demonstrably capable of resolving subcellular structure [1–4].

Despite these successes, MMF endoscopy applications are limited by knowledge of the underlying fiber’s TM because the TM changes with the fiber’s geometry: bends and twists during endoscopy change the TM, and since the distal tip at the scene cannot be accessed, the TM cannot be directly measured or calibrated using Eq. (1). Every other requirement of MMF endoscopy has been achieved experimentally: probe delivery, scene interrogation, and image formulation all follow from Eq. (1) once U is known, thus recovery of the deformed fiber’s TM is the bottleneck to general use. Proximal-only recovery methods [5] try to obtain the TM through the use of reflectors located at the distal end of the fiber. Reflection matrices within a fiber, unlike a varying and unknown scene scattering matrix outside a fiber, are constant across measurements: for K reflectors with reflection matrices $\{R_k\}$, one probe signal provides the K measurements

$$\psi_{\text{measured}}^{(k)} = U^T \tilde{R}_k U \psi_{\text{probe}} = M_k \psi_{\text{probe}}, \quad (3)$$

Where $\tilde{R}_i$ denotes the interaction with R_i as well as any prior or subsequent reflection, transmission, or propagation matrices.

This approach is appealing because it re-frames the historically physics problem as a machine learning problem: given a set of measurements $\{M_k\}$, does there exist a mapping that accurately encodes the state of an arbitrarily deformed fiber, and can that encoded state be decoded to recover the TM? Deep learning’s canonical successes never needed to ask this question because it already had empirical evidence of the answer: a finite nervous system performs vision and language, so those domains provably admit finite representations, and the architectures that succeeded there were built to exploit relevant abstract structure such as locality, invariance, and equivariance. A physics problem has no such guarantee. Whether a fiber’s transmission matrix admits a finite parameterization under deformation is a property of the fiber, and prior works provide little evidence that the inductive biases of image-domain architectures are aligned with the fiber’s dynamics. A learned model must recover the TM of a fiber state it has never encountered, and thus the model must capture and encode what the fiber does under deformation rather than what a sample of deformations looked like. In practice, the encoding must be equivariant in the informal sense that a physical action on the fiber, a bend or a twist, acts on the encoded state as a corresponding structural operation, and it is this property that allows the model to recover states it has not seen. In this work we present a new approach that suggests a finite set of graded-index (GRIN) training fibers is enough to learn that structure, and we demonstrate capabilities on a synthetic commercial grade fiber under realistic measuring conditions.

2 Related Work

Currently, the most successful approach to MMF endoscopic imaging is to make the fixed-TM assumption hold: the TM is characterized before deployment, with access to both facets, and the fiber’s geometry is then locked for the duration of imaging. The Čižmár group has carried this paradigm the furthest. They first demonstrated lensless imaging through a standard MMF [6] and then later demonstrated high-fidelity fluorescence endoscopy deep in the living brain [1].

The general TM deformation problem has also been addressed directly from the fiber side. Flaes et al. [7] showed that light transport through parabolic-profile GRIN fibers is intrinsically more robust to bending than through step-index fibers, and predicted that sufficiently precise parabolic-index fibers would make flexible probes feasible. Subsequent work pursued bending resilience through a hybrid multimode-multicore design [8], and more recently through tomographic measurement of the refractive-index perturbations of commercial GRIN fibers, showing that the measured profiles account for the observed decline of imaging under bending, identifying the fibers that endure it best, and demonstrating a flexible endoscope built on one of them [9]. This approach aims to solve the TM deformation problem by making the TM sufficiently deformation-invariant. This bend-resilience narrows the gap between the calibrated TM and the true underlying TM, but the gap still grows with the severity of the deformation and cannot be accounted for without some form of active sensing. Nevertheless, tolerance aids recovery and we argue in Section 3 that the mechanism behind the bend-resilience of graded-index fibers, low-dimensional parameterization, is what enables general TM recovery.

3 Theory

We first show that the TM for an arbitrarily deformed ideal GRIN fiber admits a finite representation while the TM for a step-index fiber does not.

3.1 Graded-Index Fiber

In the weakly-guided approximation [10], the Hamiltonian for a bent and twisted ideal GRIN fiber is the linear combination of operators (see appendix A)

$$H(z) = c_0 p^2 + c_1 r^2 + c_2 \boldsymbol{\kappa}(z) \cdot \mathbf{r} + c_3 \tau(z) L_z, \quad (4)$$

This set of deformation operators is not closed under commutation since $[p^2, r^2] = -2i(\mathbf{r} \cdot \mathbf{p} + \mathbf{p} \cdot \mathbf{r}) = -2iD$, but the set with D included does close. The non-trivial commutation relations are




Figure 1: Architecture results for 3 held-out GRIN fibers in different recovery regimes. Only every 4th column and row are shown for visual clarity.

$$\begin{aligned}
 [x_j, p_l] &= i \delta_{jl} I, & [x_j, p^2] &= 2i p_j, & [x_j, D] &= 2i x_j, & [x_j, L_z] &= -i \varepsilon_{jl} x_l, \\
 [p_j, r^2] &= -2i x_j, & [p_j, D] &= -2i p_j, & [p_j, L_z] &= -i \varepsilon_{jl} p_l, \\
 [p^2, r^2] &= -2i D, & [p^2, D] &= -4i p^2, \\
 [r^2, D] &= 4i r^2,
 \end{aligned} \tag{5}$$

where ε is the two-dimensional Levi-Civita symbol and the remaining relations $[x_j, r^2]$, $[p_j, p^2]$, $[r^2, L_z]$, $[p^2, L_z]$, and $[D, L_z]$ all vanish. Thus, $\mathfrak{g} = \text{span}\{I, x, y, p_x, p_y, L_z, p^2, r^2, D\}$ is a nine-dimensional Lie algebra that contains $H(z)$ for all deformations.

Having the general Hamiltonian contained within a finite-dimensional Lie algebra is what makes the proximal-only TM recovery problem tractable. The TM for any deformed fiber takes the path-ordered form

$$U = \mathcal{P} \exp \left(-i \int_0^L H(z) dz \right). \quad (6)$$

Dividing the fiber into segments of length Δz over which H is constant, Eq. (6) is the limit of the ordered product

$$U = \lim_{\Delta z \rightarrow 0} e^{-iH(z_M)\Delta z} \dots e^{-iH(z_2)\Delta z} e^{-iH(z_1)\Delta z}, \quad (7)$$

and neighboring segments combine through the Baker–Campbell–Hausdorff (BCH) formula,

$$e^A e^B = \exp \left(A + B + \frac{1}{2}[A, B] + \frac{1}{12}([A, [A, B]] + [B, [B, A]]) + \dots \right). \quad (8)$$

But since $\mathfrak{g}$ is closed, the expansion in Eq. (8) for the GRIN fiber can always be expressed as a linear combination of elements in the Lie algebra,

$$e^A e^B = \exp \left(\sum_i \alpha_i g_i \right), \quad g_i \in \mathfrak{g}, \quad (9)$$

and the same holds for any pair of exponentials built from $\mathfrak{g}$. An ordering change therefore adds no new terms and only changes the coefficients: applying Eq. (9) to both $e^{-A}e^{-B}$ and e^Ae^B , and then to the product of those two results gives

$$\begin{aligned} e^A e^B &= e^B e^A X, \\ X &= e^{-A} e^{-B} e^A e^B = \exp \left(\sum_i \beta_i g_i \right) \exp \left(\sum_i \gamma_i g_i \right) = \exp \left(\sum_i \delta_i g_i \right), \end{aligned} \quad (10)$$

where β , γ , and δ are each fixed by A and B through Eq. (8). Exchanging two segments therefore costs one further element of the same group rather than an expansion without end, and applying Eq. (9) across every segment of Eq. (7) collapses the ordered product to a single exponential,

$$U = f(\mathbf{a}) = \exp \left(\sum_i a_i g_i \right), \quad g_i \in \mathfrak{g}, \quad (11)$$

where $\mathbf{a}$ is fixed by the deformation profile [11, 12]. The interpretation is that while the physical geometry of any fiber is path-ordered, the TM dynamics of the GRIN fiber are not, and recovery from finite measurements is possible because those dynamics have a finite parameterization.

3.2 Step-Index Fiber

The step-index fiber replaces the quadratic well with the canonical finite-well profile, but we assume a general potential of the form $V(\mathbf{r}) = \sum_m c_m r^m$. The GRIN fiber is the special case $c_{m \neq 2} = 0$. Using the identity $[A, BC] = [A, B]C + B[A, C]$, every r^m for $m \geq 3$ in the Hamiltonian can be generated by r and r^2 , but only $m = 2$ creates a finite Lie algebra:

$$[p^2, r] = \frac{1}{r} (-2i\mathbf{r} \cdot \mathbf{p} - 1), \quad (12)$$

$$[p^2, r^4] = r^2 [p^2, r^2] + [p^2, r^2] r^2 \quad (13)$$

$$= -2i (r^2 D + D r^2). \quad (14)$$

The non-quadratic Hamiltonian is

$$H(z) = c_0 p^2 + c_1 \boldsymbol{\kappa}(z) \cdot \mathbf{r} + c_2 L_z + \sum_{m=0}^{\infty} c_{m+3} r^m \quad (15)$$

and performing the same analysis (Eqs. (6), (7), and (8)) results in a TM with no finite parameterization. The BCH formula does not terminate and the path-ordering does not become tractable. The number of operators needed to describe the fiber grows with the number of bends rather than remaining fixed, and in the continuum limit, where the curvature varies smoothly along the fiber and the number of segments is unbounded, no truncation of the expansion is uniformly valid. Recovering the TM of a step-index fiber is consequently equivalent to recovering $\boldsymbol{\kappa}(z)$ at every point along its length, which no finite set of proximal measurements can determine.

3.3 Real Fibers

For the parabolic profile, the z -dependent deformation functions $\kappa(z)$ and $\tau(z)$ never appear in the TM individually: closure consolidates each generator’s accumulated contribution into a single scalar, so the TM of Eq. (11) depends on the deformation history only through history-dependent, weighted integrals of the deformation profile, and its dynamics are exhausted by how those nine scalars evolve. A non-parabolic profile cannot be similarly compressed because each application of Eq. (8) in an effort to consolidate terms comes at the cost of an infinite expansion, and recovering the TM becomes tantamount to recovering the deformation functions themselves point by point along the fiber.

In practice, commercial GRIN fibers are not perfectly parabolic, and contain many defects that break this ideal structure. Fortunately, the defects in the spatial refractive profile for GRIN fibers are perturbative, and parameter consolidation introduces higher-order BCH terms that decay rapidly, and general recovery from finite measurements only needs the dynamics to be well approximated by some finite representation.

Demonstrating general recovery requires data sampled from fibers whose curvature admits no convenient parameterization and this data would be expensive to obtain experimentally. Consequently, we establish the architecture’s capabilities on synthesized fibers, where deformation states are sampled from a continuum. Details about the Hamiltonian used to describe the fiber’s dynamics can be found in Appendix A, and details about how the fiber is deformed and how the TM and reflection measurements are obtained can be found in Appendix B.

4 Methods

4.1 Proximal-Only Recovery System for MMF Endoscopy

In a physically realizable MMF endoscopic system, an input signal propagates down the fiber, interacts with the reflectors and the scene, and their respective return signals are measured back at the proximal end. The $N \times N$ reflection and scene matrices can be obtained from N linearly independent input signals using an amplitude + phase measuring technique such as off-axis holography.




Figure 2: Visual representation of our proximal-end recovery system. The input signal (blue) propagates down the fiber, reflects off of reflectors R_1 , R_2 , R_3 , and the scene S , and then propagates back. The TM is determined from the measurement signals (red) and then uses the scene signal (green) to construct the final image.

All matrices and signals are represented in the Laguerre-Gauss spatial mode basis — the propagation-invariant modes (PIM) for the ideal GRIN fiber. These PIMs carry a complex $e^{i\ell\varphi}$ factor, where propagation back to the proximal end after reflection is the transpose composed with the $\ell \rightarrow -\ell$ operator P for circularly polarized light. Writing $\mathcal{P}[X] \equiv PXP$, the round trip reflection measurements are $M_i = \mathcal{P}[U^T] R_i U$.

Since these measurements are bilinear, at least three reflectors are needed to break all symmetries and uniquely determine U [5, 13]. In this work we use a straight partial reflector R_1 , a tilted partial reflector R_2 , and a second tilted partial reflector R_3 with its tilt axis making a $\pi/4$ angle with R_2 ’s tilt axis. We note that neither R_i nor $\hat{R}_i$ need to be explicitly known or fabricated with extreme precision in practice. The architecture only requires that all bilinear-induced symmetries are sufficiently broken by the set of reflectors.

The architecture solves for a fiber’s bent TM, U_2 , given the fiber’s unbent TM, U_1 , and the three reflector measurements $\{M_1, M_2, M_3\}$. The unknown U_2 is obtained indirectly by solving for $V = U_2 U_1^\dagger$ instead, where $U_2 = V U_1$ and V is

the operator that describes the mode couplings due to the deformation. Using $\mathcal{P}[XY] = \mathcal{P}[X]\mathcal{P}[Y]$,

$$M_i = \mathcal{P}[(VU_1)^T] R_i (VU_1) = \mathcal{P}[U_1^T] \mathcal{P}[V^T] R_i V U_1 = \mathcal{P}[U_1^T] A_i U_1 \quad (16)$$

$$A_i = \mathcal{P}[V^T] R_i V = \mathcal{P}[U_1^*] M_i U_1^\dagger, \quad (17)$$

where A_i is a known since both M_i and U_1 are known.

4.2 Architecture

We use a bilinear architecture [14, 15] with one encoder per measurement A_i , and replace the correspondence decoder in [15] with a MLP that outputs V :

$$h_{jk}^i = \langle \phi_j^i | A_i | \psi_k^i \rangle, \quad (18)$$

$$g_n^i = \sigma(W_{jkn} h_{jk}^i + c_n), \quad (19)$$

$$\text{cat}[\text{Re}(V), \text{Im}(V)] = \text{MLP}(\text{cat}[\mathbf{g}^1, \mathbf{g}^2, \mathbf{g}^3]), \quad (20)$$

where σ is a non-linear activation function (we use tanh in this paper) and the MLP contains two hidden layers with the leaky ReLU activation function ($\mathbf{g} \rightarrow \mathbf{z}_1 \rightarrow \mathbf{z}_2$). Outputting the entire V from a single MLP is impractical for large number of modes, and the presented architectures instead partition V into subsets and then use a separate, smaller MLP for each subset, $V_i = \text{MLP}_i(\mathbf{z}_2)$. Complex $\mathbf{h}$ values are handled in σ by $\mathbf{h} \rightarrow \text{cat}[\|\mathbf{h}\|, \text{Re}(\mathbf{h}), \text{Im}(\mathbf{h})]$. No learned or prescribed residual gating mechanism is used in the presented architectures for simplicity, but the iterative refinement procedure still applies: an initial output $\hat{V}_1$ can be applied to U_1 within A_i to "deform" U_1 toward U_2 . Inference is then performed using the updated A_i matrices, and the estimate composes as $V = \hat{V}_2 \hat{V}_1$, $\hat{U}_2 = V U_1$. Two inference passes are applied during training and testing.

In practice, illumination and detection involve distinct optical systems and each imposes its own transformation on the measurement, so what the architecture sees is BA_iC rather than A_i , where the transformations B and C are fixed properties of the system rather than of the fiber [16]. These transformations break the reciprocity relation of Eq. (1) but they do not obstruct recovery since each feature h_{jk}^i is a bilinear form in the measurements and the filters are mutable: for invertible B, C and synthetically trained filters ψ, ϕ , the experimentally trained (or fine-tuned) filters would become $(B^{-1})^\dagger \phi$ and $C^{-1} \psi$, leaving h_{jk}^i intact (experimental filters for non-invertible transformations would instead move towards the available orthogonal projection, h_{jk}^i would change, and downstream weights would likewise need to be fine-tuned).

More generally speaking, the encoder does not commit to the basis in which the measurement is expressed: A_i is an operator and its rank is a basis-free property, thus the set of features the encoder can compute is unchanged under any invertible change of basis on either side of the measurement. This distinguishes it from the locality and translation structure that image-domain architectures impose on a pixel basis that has no counterpart in a measurement whose content is mode coupling.

5 Results

We trained our architecture under the protocol described in Table 2 on a synthesized DRAKA WideCap OM5 50 μm core fiber at $\lambda = 532$ nm with both circular-polarization states (σ_+, σ_-) and $N = 300$ modes per polarization (600 total modes). We trained three separate models with varying levels of zero-mean gaussian noise, $\sigma = \{0.0, 0.04, 0.16\}$, added to $\{M_i\}$. We also trained a model on a synthesized step-index fiber as a reference with $N = 303$ using the Linearly Polarized basis and zero-mean $\sigma = 0.02$ gaussian noise added to its $\{M_i\}$. The evaluation metric is fiber fidelity $F = |\text{tr}(V_{\text{pred}}^\dagger V_{\text{gt}})| / (\|V_{\text{pred}}\|_F \|V_{\text{gt}}\|_F)$ where V_{pred} is the model output and V_{gt} is the ground truth matrix.

5.1 Learned Curvature Encoding

To examine how the architecture organizes measurement information from these arbitrarily deformed GRIN fibers, we constructed 820 fibers each wound into a single circle of constant curvature, with radii of curvature between 3 and 60 cm. While these fibers are unphysical (many of them wind into themselves), they isolate a single scalar deformation parameter and allow it to be swept continuously. Furthermore, they appear in neither the training nor the testing set, and radii of curvature larger than 40 cm are never seen by the trained models. Figure 4 shows the main MLP's second hidden layer activations $\mathbf{z}_2$ for each fiber projected onto the first three principal components. The architecture only sees the three measurements $\{M_i\}$, and plotting these projections as a function of curvature shows the parameter traversing




Figure 3: Recovery fidelity on the thousand held-out fibers for the three GRIN models trained at detector noise $\sigma = 0, 0.04$ and 0.16 , and the step-index model. each evaluated at its own σ . Top: Scatter plot of F as a function of the minimum bend radius, with a rolling mean over 80 fibers. Individual scatter points are shown for the $\sigma = 0.16$ model. Bottom: Histogram of the same fidelities.

Table 1: Fidelity recovery for the $\sigma = 0.16$ GRIN model on held-out testing data with mean initial fidelity $F = |\text{tr}(U_2 U_1^\dagger)|/(\sqrt{N} \|U_2 U_1^\dagger\|_F) = 0.1816$.

Subset	n	Mean F	p_5	Worst
All held-out fibers	1000	0.897	0.551	0.045
Min bend radius ≥ 5 cm	808	0.953	0.869	0.603
Min bend radius ≥ 7 cm	692	0.964	0.919	0.671
Min bend radius ≥ 10 cm	566	0.970	0.934	0.671

a three-dimensional manifold within the model’s latent space. The first two components correspond to a local phase and the third component corresponds to a coarse global structure. The encoded curvature loses its structure around $R = 7$ cm, in agreement with the testing dataset results, and provides one possible explanation why some stray fibers with large R in Figure 3 were poorly recovered: their latent vector sits near the region with broken structure and was incorrectly decoded. The step-index model, despite the exact same encoder, could not encode the curvature as well, which is in agreement with [7].

Table 2: Training protocol

Parameter	Value
Training deformations	19,000
Testing deformations	1000
Loss ¹	$\min_{\theta} \ Q(\theta)V_{\text{pred}} - V\ _F^2 / \ V\ _F^2$
Optimizer	AdamW
Learning rate	3×10^{-4}
Schedule	cosine annealing to 0
Epochs	40
Batch size	24
Inference passes	2

¹ In the circular polarization basis $Q(\theta) = \text{Diag}(e^{i\theta} I_N, e^{-i\theta} I_N)$, where θ is determined analytically.

We then looked to see if curvature is encoded by the model in higher order harmonics within $\mathbf{z}_2$ with a dense sampling of the same constant curvature fibers from $R = 40$ cm to $R = 20$ cm. Figure 5 shows that curvature is indeed very clearly encoded in higher order harmonics, and extends until the 12-14th harmonic for this range of R .

5.2 Learned Relative Positional Encoding

Another encoded property is the location of a particular deformation along the fiber. For a fiber with a localized defect that is partitioned into n_{total} sections, the TM can be represented by

$$U(n) = (e^{-iH_0\Delta z})^{m-n} U_{\text{defect}} (e^{-iH_0\Delta z})^n, \quad (21)$$

where there are m straight fiber segments, the defect is described by the remaining $n_{\text{total}} - m$ segments, n indicates where the defect exists along the fiber, and H_0 is the unbent Hamiltonian. Changing n corresponds to changing X in Eq. (10), and to see if the model learned to encode this relational structure, we constructed 10,800 fibers in which the same localized spiral deformation (complex Gabor wavelet) is translated along the middle quarter of the fiber (Figure 6a). Moving the defect conjugates the TM by straight propagation, $U(n) = e^{-iH_0n\Delta z} U(0) e^{iH_0n\Delta z}$, so each element only acquires a phase $e^{-i(\beta_j - \beta_k)n\Delta z}$, where β_j are the propagation constants, and position enters the TM in the same way it enters a rotary positional encoding.

Figure 6b shows that the architecture learned this structure. The latent $\mathbf{z}_2$ oscillates with a period within 0.4% of $2\pi/\Omega$, and over a few millimeters its two largest principal components trace a circle, so the defect position is encoded as an angle while a slower component across the full sweep carries the coarse position along the fiber. Neither the defect position nor the self-imaging period is supplied to the architecture at any point, and like the curvature ablation, this relational structure is something it arrives at from the measurements alone.

6 Discussion and Limitations

The parameters of the simulated fiber and proximal-only system were chosen to go beyond proof-of-concept: the refractive index perturbations are the measured profiles of a Prysmian DRAKA WideCap OM5 50 μm core fiber and the Hamiltonian of Eq. (35) contains the spin-orbit coupling together with all three birefringence contributions. The deformation regime spans minimum bend radii from 3 to 40 cm with out-of-plane displacement, non-adiabatic curvature changes, and twists. The straight calibration U_1 retains a fidelity of 0.12-0.22 (mean 0.18) against the deformed U_2 , indicating that these are deformations under which a fixed calibration is unusable. The measurements are collected in a realistic, non-ideal manner where the fiber is perturbed independently during each probe signal (Eq. (36)) so that no M_i corresponds to a single fiber state, and with system noise added on top of that (Figure 8). Each of these choices moves the setting toward a working instrument and away from the regime in which recovery would be easiest.

Furthermore, our construction poses the hardest version of the recovery problem. Each deformation operator V is recovered in a single step between the straight calibration state U_1 and a severely deformed U_2 , and nothing in the formulation requires this, since $V = U_2 U_1^\dagger$ holds for any previously known state in place of U_1 . A deployed endoscope deforms continuously, so the most recently recovered TM could and should serve as the reference where each update only needs to recover the incremental deformation accumulated over one time step, which is a far better conditioned




Figure 4: The three largest principal components of $\mathbf{z}_2$ for the initial 620 wound fibers plus an additional 200 fibers in the $R \in [40, 60]$ cm range for the $\sigma = 0.16$ GRIN fiber (top) and step-index fiber (bottom), colored by radius of curvature. Gray lines indicate trajectory path. The architectures never saw bends with $R \geq 40$ cm, so that curvature regime is outside of the training distribution but still traverses the manifold of both architectures as if it was in the distribution and demonstrates generalized capabilities. However, curvature information is much better compressed in the GRIN encoder, maintaining low-dimensional structure until about $R = 7$ cm, whereas low-dimensional representation in the step-index encoder breaks down near $R = 20$ cm.





Figure 5: Curvature of the $\sigma = 0.16$ GRIN model encoded in $\mathbf{z}_2$ as higher-order harmonics.

problem than the full excursion from straight. From this perspective, we consider the fidelities in Table 1 to be a lower bound on what the same architecture would achieve operating as a closed tracking loop with appropriate algorithms.


While the fibers are synthetic, the path from these results to an experimentally verified device does not require discarding the synthetic data. Every term in Eq. (35) corresponds to an identified physical effect or fiber defect rather than a fitted degree of freedom: the parabolic potential and its measured perturbation, the bend-induced tilt, the twist rotation, the spin-orbit coupling, and the three birefringence contributions each have a separate physical origin, so a discrepancy between simulated and experimental measurements can be attributed to particular terms and corrected rather than absorbed into a global mismatch. The same holds for the measurement chain, where the sources we do not currently model are characterizable on a bench, and once characterized they enter the simulator in the same way that the acquisition jitter and system noise already do. Figure 3 shows that recovery is robust to measurement noise, with the $\sigma = 0.16$ measurements containing significant magnitude and phase errors, indicating that the architecture is not relying on artifacts particular to noise-free data. We therefore expect a practical route to an instrument to be pretraining on synthetic data made as faithful as the characterization allows, followed by fine-tuning on a comparatively small experimental set.

7 Conclusion

In this work we introduce a bilinear neural network architecture for proximal-only TM recovery that achieves an average fidelity recovery of $F = 0.897$ on arbitrarily deformed GRIN MMFs with 600 guided modes ($F \geq 0.95$ when restricted to bending curvatures larger than 5 cm). We also perform several ablations demonstrating the learned model’s ability to identify structure and generalize to fibers outside of the training distribution. The next step is to reproduce these capabilities on experimentally acquired measurements, and we believe the results here establish that the effort is warranted.



(a) The first ($z_0 = 375$ mm) and last ($z_0 = 625$ mm) of the 10,800 fibers, where z_0 (indicated by the dot) is the center of the spiral deformation.



(b) The three largest principal components of $\mathbf{z}_2$ for all 10,800 fibers, colored by z_0 .

Figure 6: A localized spiral deformation translated along the middle quarter of the fiber. None of these fibers appear in the training or testing sets.

References

- [1] Sergey Turtaev, Ivo T. Leite, Tristan Altwegg-Boussac, Janelle M. P. Pakan, Nathalie L. Rochefort, and Tomáš Čížmár. High-fidelity multimode fibre-based endoscopy for deep brain in vivo imaging. *Light: Science & Applications*, 7:92, 2018. doi: 10.1038/s41377-018-0094-x.
- [2] Sebastian A. Vasquez-Lopez, Raphaël Turcotte, Vadim Koren, Martin Ploschner, Zahid Padamsey, Martin J. Booth, Tomáš Čížmár, and Nigel J. Emptage. Subcellular spatial resolution achieved for deep-brain imaging in vivo using a minimally invasive multimode fiber. *Light: Science & Applications*, 7(1):110, 2018. doi:

- 10.1038/s41377-018-0111-0.
- [3] Shay Ohayon, Antonio M. Caravaca-Aguirre, Rafael Piestun, and James J. DiCarlo. Minimally invasive multimode optical fiber microendoscope for deep brain fluorescence imaging. *Biomedical Optics Express*, 9(4):1492–1509, 2018. doi: 10.1364/BOE.9.001492.
 - [4] Zhong Wen, Zhenyu Dong, Qilin Deng, Chenlei Pang, Clemens F. Kaminski, Xiaorong Xu, Huihui Yan, Liqiang Wang, Songguo Liu, Jianbin Tang, Wei Chen, Xu Liu, and Qing Yang. Single multimode fibre for in vivo light-field-encoded endoscopic imaging. *Nature Photonics*, 17(8):679–687, 2023. doi: 10.1038/s41566-023-01240-x.
 - [5] R. Y. Gu, Reza Nasiri Mahalati, and Joseph M. Kahn. Design of flexible multi-mode fiber endoscope. *Optics Express*, 23(21):26905–26918, 2015. doi: 10.1364/OE.23.026905.
 - [6] Tomáš Čižmár and Kishan Dholakia. Exploiting multimode waveguides for pure fibre-based imaging. *Nature Communications*, 3:1027, 2012. doi: 10.1038/ncomms2024.
 - [7] Dirk E. Boonzajer Flaes, Jan Stopka, Sergey Turtaev, Johannes F. de Boer, Tomáš Tyc, and Tomáš Čižmár. Robustness of light-transport processes to bending deformations in graded-index multimode waveguides. *Physical Review Letters*, 120:233901, 2018. doi: 10.1103/PhysRevLett.120.233901.
 - [8] Yang Du, Sergey Turtaev, Ivo T. Leite, Adrian Lorenz, Jens Kobelke, Katrin Wondraczek, and Tomáš Čižmár. Hybrid multimode-multicore fibre based holographic endoscope for deep-tissue neurophotonics. *Light: Advanced Manufacturing*, 3(3):408–416, 2022. doi: 10.37188/lam.2022.029.
 - [9] Sergey Turtaev, Tomáš Tyc, Ulf Poßner, Tina Eschrich, Torsten Poßner, Yang Du, André Gomes, Bernhard Messerschmidt, and Tomáš Čižmár. Deformation enduring conveyance of structured light through multimode waveguides and its exploitation for flexible hair-thin endoscopes, 2025.
 - [10] Allan W. Snyder and John D. Love. *Optical Waveguide Theory*. Chapman and Hall, London, 1983.
 - [11] Wilhelm Magnus. On the exponential solution of differential equations for a linear operator. *Communications on Pure and Applied Mathematics*, 7(4):649–673, 1954. doi: 10.1002/cpa.3160070404.
 - [12] James Wei and Edward Norman. Lie algebraic solution of linear differential equations. *Journal of Mathematical Physics*, 4(4):575–577, 1963. doi: 10.1063/1.1703993.
 - [13] Szu-Yu Lee, Vicente J. Parot, Brett E. Bouma, and Martin Villiger. Reciprocity-induced symmetry in the round-trip transmission through complex systems. *APL Photonics*, 5(10):106104, 2020. doi: 10.1063/5.0021285.
 - [14] Roland Memisevic and Geoffrey Hinton. Unsupervised learning of image transformations. In *2007 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8. IEEE, 2007. doi: 10.1109/CVPR.2007.383036.
 - [15] Cole Reynolds. Neural phase correlation, 2026. URL <https://arxiv.org/abs/2606.18496>.
 - [16] Shreyas Bharadwaj, Gyeong Hun Kim, Brett E. Bouma, and Martin Villiger. Matrix symmetry and normality in multimode fiber transport. *APL Photonics*, 11(9):096101, 2026. doi: 10.1063/5.0337399.
 - [17] Moty Heiblum and Jay H. Harris. Analysis of curved optical waveguides by conformal transformation. *IEEE Journal of Quantum Electronics*, 11(2):75–83, 1975. doi: 10.1109/JQE.1975.1068563.
 - [18] Martin Plöschner, Tomáš Tyc, and Tomáš Čižmár. Seeing through chaos in multimode fibres. *Nature Photonics*, 9(8):529–535, 2015. doi: 10.1038/nphoton.2015.112.
 - [19] Hui Cao, Tomáš Čižmár, Sergey Turtaev, Tomáš Tyc, and Stefan Rotter. Controlling light propagation in multimode fibers for imaging, spectroscopy, and beyond. *Advances in Optics and Photonics*, 15(2):524–612, 2023. doi: 10.1364/AOP.484298.
 - [20] R. Clark Jones. A new calculus for the treatment of optical systems. i. description and discussion of the calculus. *Journal of the Optical Society of America*, 31(7):488–493, 1941. doi: 10.1364/JOSA.31.000488.
 - [21] Ivan P. Kaminow. Polarization in optical fibers. *IEEE Journal of Quantum Electronics*, 17(1):15–22, 1981. doi: 10.1109/JQE.1981.1070626.
 - [22] R. Ulrich, S. C. Rashleigh, and W. Eickhoff. Bending-induced birefringence in single-mode fibers. *Optics Letters*, 5(6):273–275, 1980. doi: 10.1364/OL.5.000273.
 - [23] R. Ulrich and A. Simon. Polarization optics of twisted single-mode fibers. *Applied Optics*, 18(13):2241–2251, 1979. doi: 10.1364/AO.18.002241.

A Propagation Hamiltonian

A.1 Scalar Approach

The scalar field in a weakly guiding fiber obeys the Helmholtz equation $\nabla^2 E + k_0^2 n^2(x, y, z) E = 0$. The ansatz $E = \psi(x, y, z) e^{ikz}$ with $k = n_0 k_0$, the free-space wavenumber $k_0 = 2\pi/\lambda$ and n_0 the index on the fiber axis, gives

$$\nabla^2 \psi + 2ik \partial_z \psi + (k_0^2 n^2 - k^2) \psi = 0. \quad (22)$$

When ψ varies slowly w.r.t z on the scale of a wavelength, the approximation $|\partial_z^2 \psi| \ll |2k \partial_z \psi|$ can be made and the $\partial_z^2 \psi$ term can be neglected. What remains is a 2-D Schrödinger-like equation on the xy -plane in which the propagation coordinate z plays the role of time,

$$i \partial_z \psi = H \psi, \quad H = \frac{p^2}{2k} + V(\mathbf{r}), \quad V = \frac{k}{2} \left(1 - \frac{n^2}{n_0^2} \right), \quad (23)$$

with $\mathbf{p} = -i\nabla_\perp$ acting on the two transverse (x, y) coordinates only. For the ideal parabolic profile $n^2(r) = n_0^2 [1 - 2\Delta (r/a)^2]$, with core radius a and

$$\Delta = \frac{n_0^2 - n_{\text{cl}}^2}{2n_0^2} = \frac{\text{NA}^2}{2n_0^2} \quad (24)$$

is the relative index contrast between the axis and the cladding, Eq. (23) gives a harmonic potential,

$$V_0 = \frac{k}{2} \Omega^2 r^2, \quad \Omega = \frac{\sqrt{2\Delta}}{a}, \quad (25)$$

and the unperturbed problem is the 2-D isotropic oscillator, whose eigenmodes are the Laguerre–Gauss modes $|\ell, m\rangle$ with propagation constants β_g for each mode group index $g = |\ell| + 2m$.

Bending enters through the geometry. A fiber bent to a local radius of curvature R is equivalent to a straight fiber whose index profile is tilted across the core [7, 17]. To first order in r/R the equivalent profile is $n_{\text{eq}}^2 \approx n^2 (1 + 2\boldsymbol{\kappa} \cdot \mathbf{r})$ where $\boldsymbol{\kappa}$ is the curvature vector with magnitude $1/R$ pointing along the bend normal, and substituting into Eq. (23) adds the term $-k \boldsymbol{\kappa} \cdot \mathbf{r}$, linear in the transverse coordinate. The bend also compresses the glass on the inside and stretches it on the outside, changing the density and with it the index. Writing $n - 1$ as proportional to density and using $\varepsilon_{xx} = \varepsilon_{yy} = -\sigma \varepsilon_{zz}$, where σ is the Poisson ratio, that correction leaves the form of the term alone and rescales the curvature itself [18],

$$\eta = 1 - (1 - 2\sigma) \frac{n_0 - 1}{n_0}, \quad (26)$$

so the fiber bends as though its curvature were $\eta \boldsymbol{\kappa}$. For fused silica $\sigma = 0.17$, giving $\eta = 0.786$ at $n_0 = 1.4783$, against the 0.77 ± 0.02 measured in [18]. The bend term is therefore

$$V_{\text{bend}}(z) = -k \eta \boldsymbol{\kappa}(z) \cdot \mathbf{r}. \quad (27)$$

Twist is a cross-section that rotates along the fiber [10], through an accumulated angle $\theta(z)$ at rate $\tau(z) = d\theta/dz$, and collecting all contributions gives the ideal-fiber Hamiltonian

$$H_{\text{ideal}}(z) = \frac{p^2}{2k} + \frac{k\Omega^2}{2} r^2 - k \eta \boldsymbol{\kappa}(z) \cdot \mathbf{r} - \tau(z) L_z. \quad (28)$$

Realistically, a manufactured fiber is not an ideal parabola. Its measured index deviates from the parabolic profile by the perturbation $\delta n^2(x, y)$, and substituting $n^2 = n_0^2 [1 - 2\Delta (r/a)^2] + \delta n^2$ into Eq. (23) leaves the harmonic potential untouched and adds the perturbative term,

$$\delta V(\mathbf{r}) = -\frac{k}{2n_0^2} \delta n^2(x, y). \quad (29)$$

Crucially, this term breaks the degeneracy of the mode groups: the ideal parabola assigns a single β_g to every state of a given g , and δV non-trivially perturbs each mode's propagation constant from the group value.

A.2 Vector Approach

GRIN fibers are not currently manufactured well enough for the scalar, single-polarization approach to be valid [9, 19]: ellipticity, refraction perturbations, and stresses introduced during the fiber draw couple $\mathbf{E}$'s polarization states and should be accounted for. The E-field obeys the vector wave equation [10]

$$\nabla^2 \mathbf{E} + k_0^2 n^2 \mathbf{E} + \nabla (\mathbf{E} \cdot \nabla \ln n^2) = 0. \quad (30)$$

Using the two-state Jones vector ψ [20] in the E-field ansatz, the scalar Hamiltonian in Eq. (23) carries over unchanged but the last term of Eq. (30) introduces the spin-orbit coupling term H_{SO} ,

$$H_{\text{SO}} = -\frac{1}{4k} \begin{pmatrix} \partial_- w_+ & \partial_- w_- \\ \partial_+ w_+ & \partial_+ w_- \end{pmatrix}, \quad \partial_{\pm} = \partial_x \pm i\partial_y, \quad w_{\pm} = \partial_{\pm} \ln n^2, \quad (31)$$

where ψ is projected onto $\mathbf{e}_{\pm} = (\mathbf{e}_x \pm i\mathbf{e}_y)/\sqrt{2}$ because the architecture assumes circular polarization.

Material anisotropy contributes three further terms acting on the polarization index. (1) Frozen stresses leave a fixed linear birefringence [21] whose slow axis sits at an angle α set by the quadrupole component of the measured profile:

$$H_1 = \frac{b_1}{2} (e^{-2i\alpha} \sigma_+ + e^{2i\alpha} \sigma_-), \quad (32)$$

where $\sigma_{\pm} = (\sigma_x \pm i\sigma_y)/2$. (2) Bending compresses the glass on the inside of the bend and stretches it on the outside. The resulting stress-optics birefringence [22] has its principal axes set by the bend plane and its magnitude is quadratic in curvature. Define $\phi(z)$ to be the angle $\kappa(z)$ makes with the x -axis of the co-rotating spatial profile,

$$H_2 = \frac{b_2 \kappa^2}{2} (e^{-2i\phi} \sigma_+ + e^{2i\phi} \sigma_-), \quad (33)$$

where $\kappa = |\kappa|$. This is identical to H_1 with the bend plane in place of the frozen axis and $b_2 \kappa^2$ in place of b_1 with the sign of b_2 fixing which of the two axes is slow. (3) Twisting shears the glass, and the photoelastic response to that shear is a circular birefringence

$$H_3 = b_3 \tau(z) \sigma_z. \quad (34)$$

Collecting every contribution gives the Hamiltonian used for our synthesized fibers,

$$H(z) = \frac{p^2}{2k} + \frac{k\Omega^2}{2} r^2 + \delta V(\mathbf{r}) - k \eta \kappa(z) \cdot \mathbf{r} - \tau(z) L_z + H_{\text{SO}} + H_1 + H_2 + H_3. \quad (35)$$

B Fiber Model and Dataset Synthesis

B.1 Index profile and mode basis

We simulate the Prysmian DRAKA WideCap OM5 fiber, with a 50 μm core and NA 0.200, at $\lambda = 532$ nm. The index profile is the ideal parabolic profile with the measured perturbation $\Delta n^2(x, y)$ reported in [9].

B.2 Deformation synthesis

Each deformation is generated by partitioning the fiber into n equal-length sections, where n is uniformly sampled from $\{1, 2, 3, 4, 5, 6\}$. Each section is assigned a curvature vector κ whose magnitude corresponds to a radius of curvature drawn uniformly from 3 to 40 cm and whose direction is drawn uniformly on the circle, so that the deformation is not confined to a single plane. The resulting piecewise-constant $\kappa(z)$ is convolved with a Gaussian kernel to produce a continuous curvature profile, which makes the curvature vary smoothly along the fiber while remaining non-adiabatic. A twist $\tau(z)$ is applied over the full length, with a total accumulation uniformly sampled from $[-2 \text{ rad}, 2 \text{ rad}]$.

B.3 Transmission matrix construction

The smoothed deformation profile is discretized into 100 segments of constant κ and τ , the Hamiltonian in Eq. (35) is evaluated on each segment in the retained mode basis, and the transmission matrix is the ordered product of segment propagators as in Eq. (7). The undeformed transmission matrix U_1 is obtained from the same construction with $\kappa = 0$ and $\tau = 0$, so that U_1 carries the fiber's defects.

Table 3: Simulated DRAKA fiber parameters.

Symbol	Meaning	Value	Source
<i>Fiber (Prysmian DRAKA WideCap OM5)</i>			
λ	build wavelength	532 nm	N.A.
n_0	index on axis	1.4783	manufacturer
NA	numerical aperture	0.200	manufacturer
a	core radius	25 μm	manufacturer
R_{cl}	cladding radius	62.5 μm	manufacturer
Δ	index contrast	9.15×10^{-3}	Eq. (24)
L	fiber length	1.0 m	N.A.
k_0	vacuum wavenumber	$1.181 \times 10^7 \text{ m}^{-1}$	$2\pi/\lambda$
k	axial wavenumber	$1.746 \times 10^7 \text{ m}^{-1}$	$n_0 k_0$
Ω	oscillator frequency	5412 m^{-1}	$\sqrt{2\Delta}/a$
σ	Poisson ratio	0.17	fused silica [18]
η	elasto-optic curvature factor	0.786	Eq. (26)
<i>Polarization coefficients</i>			
α	frozen-stress axis	$+18.5^\circ$	[9], quadrupole of δn^2
b_1	frozen linear biref.	3.0 rad m^{-1}	no source
b_2	bend stress-optic	$6.25 \times 10^{-3} \text{ rad m}$	[22]
b_3	twist activity	0.14	[23]; 0.13–0.16

Table 4: Step-index fiber parameters. The fiber is mode-matched to the DRAKA fiber of Table 3: same core radius, same n_0 , same wavelength, same defect field δn^2 , same polarization coefficients, and similar number of guided modes.

Symbol	Meaning	Value	Source
<i>Fiber (step-index, mode-matched to the DRAKA fiber)</i>			
λ	build wavelength	532 nm	matched
n_0	core index	1.4783	matched
n_{cl}	cladding index	1.47131	Eq. (24)
NA	numerical aperture	0.1436	chosen for 452 guided modes
a	core radius	25 μm	matched
R_{cl}	cladding radius	62.5 μm	matched
Δ	index contrast	4.72×10^{-3}	Eq. (24)
V	normalized frequency	42.40	$k_0 a \text{ NA}$
L	fiber length	1.0 m	matched
k_0	vacuum wavenumber	$1.181 \times 10^7 \text{ m}^{-1}$	$2\pi/\lambda$
k	axial wavenumber	$1.746 \times 10^7 \text{ m}^{-1}$	$n_0 k_0$
Ω_{eff}	oscillator frequency	3048 m^{-1}	$\sqrt{2\Delta}/a$
σ	Poisson ratio	0.17	fused silica [18]
η	elasto-optic curvature factor	0.786	Eq. (26)
<i>Polarization coefficients (unchanged from Table 3)</i>			
α	frozen-stress axis	$+18.5^\circ$	matched
b_1	frozen linear biref.	3.0 rad m^{-1}	matched
b_2	bend stress-optic	$6.25 \times 10^{-3} \text{ rad m}$	matched
b_3	twist activity	0.14	matched

B.4 Reflector stack and measurement acquisition

The distal stack consists of a straight partial reflector $R_1 = \rho_1 I$, a tilted partial reflector R_2 , and a second tilted partial reflector R_3 whose tilt axis makes an angle of $\pi/4$ with that of R_2 . Both R_2 and R_3 are tilted 0.6 deg. Reflector R_2 is simulated 3mm behind R_1 , and R_3 is simulated 5mm behind R_2 . The transmission matrices T_i are the scalars $t_i I$ that satisfy energy conservation.

Each measurement M_i is assembled from N linearly independent probe signals, and each probe is subject to an independent rigid misalignment of the launch. We model it to first order in the four generators $J_a \in \{X, Y, P_x, P_y\}$ of




Figure 7: Generated fiber deformations in the held-out test set.

transverse displacement and tilt, so that the measured matrix becomes

$$M' = M + i \sum_a \varepsilon_a J_a M D_a + i M \sum_a J_a D_a, \quad D_a = \text{diag}(d_{j,a}), \quad (36)$$

where $d_{j,a}$ is a randomly sampled from a Gaussian independently for every probe j , whose standard deviation σ_j is itself assigned per fiber, stratified over $\sigma_j \in [0, 0.10]$ so that the population spans the range with mean 0.05, and $\varepsilon_a = (1, 1, -1, -1)$ is the parity of each generator under the reciprocity flip P . The σ_j varies for each synthetic fiber's $\{M_i\}$ measurements, but U_1 is constructed with a smaller $\sigma_j = 0.01$ since care can be taken before training to ensure it's more accurate with techniques like multi-look averaging. The measurements are assumed single-look and therefore contain more "jitter" noise. Two terms appear because the same physical misalignment is seen on the launch and again on the return path, and the right-acting operator $\sum_a J_a D_a$ is dense rather than diagonal. Zero-mean Gaussian noise of standard deviation σ is then added to the reflection measurements.




Figure 8: The third reflector measurement M_3 of a single deformed fiber at three noise levels. The top row is the measurement as the architecture receives it, where hue is phase and brightness is magnitude. The bottom row is the phase deviation of each matrix element from its ideal value, $\arg(M_{\text{ideal}}^* \circ M_{\text{measured}})$, where $\circ$ is the elementwise product and M_{ideal}^* is the measurement that a single, unperturbed fiber state would produce without added noise. The $\sigma = 0$ column therefore isolates the per-probe deformation applied during acquisition, which is present in every column.