Supplementary Materials

Deep Learning Driven Enhancement of Optical Vortex Line Robustness in Atmospheric Turbulence

Dmitrii Tsvetkov*†, Danilo Gomes Pires†, and Natalia Litchinitser

Department of Electrical and Computer Engineering, Duke University, Durham, North Carolina 27708, USA

**dmitrii.tsvetkov@duke.edu*

†These authors contributed equally to this work

S1. Alphabet-11 Optical Knots

In this section, we define each knot in Alphabet-11, a set of 11 different optical knots comprising nine Hopf links with varying geometries (H1–H9) and two trefoil knots (T1 and T2). Figure S1 illustrates these knots, showing their 3D structures, corresponding field distributions in the central plane, and their respective Laguerre-Gaussian (LG) spectra.

The LG spectrum coefficients, denoted as $c_{l,p}$, are obtained by decomposing the electric field of each knot, E_{knot}, as follows:

$$E_{\text{knot}} = \sum_{l,p} c_{l,p}\, \text{LG}_{l,p}(R, \varphi, z), \tag{S1}$$

where LG modes are defined as:

$$\text{LG}_{p,l}(R, \phi, z) = \sqrt{\frac{p!}{\pi(|l|+p)!}} \frac{R^{|l|} e^{il\phi}}{w_0^{|l|+1}} \frac{(1 - iz/z_R)^p}{(1 + iz/z_R)^{p+|l|+1}} e^{-\frac{R^2}{2w_0^2(1+iz/z_R)}} L_p^{|l|}\left(\frac{R^2}{w_0^2(1+z^2/z_R^2)}\right). \tag{S2}$$

Here, l is the azimuthal index, p is the radial index, $z_{\text{R}} = k_0 w_0^2$ is the Raleigh range, k_0 is the wavenumber, L_p^l are the generalized Laguerre polynomials, and w_0 is the beam waist parameter.

Below, we provide a detailed description of each knot, outlining its construction and defining properties, with all weights $c_{l,p}$ listed in Table S1.

H1, H2, H3 Hopf Links

The H1, H2, and H3 Hopf links are classical isolated optical knots that can all be obtained from the Milnor map [1]:

$$q_{\text{H1}} = (u - v)(u + v), \tag{S3}$$

where u and v are defined by the inverse stereographic projection, using the complex coordinates for the 3-sphere:

$$u(R, \phi, z) = \frac{R^2 + z^2 - 1 + 2iz}{R^2 + z^2 + 1}, \qquad v(R, \phi, z) = \frac{2Re^{i\phi}}{R^2 + z^2 + 1}. \tag{S4}$$

In cylindrical coordinates at the $z = 0$ plane, after removing the common divisor, this transforms into:

$$q_{\text{H1,H2,H3}}(R, \varphi, z = 0) = 1 - 2R^2 - 4R^2 \exp 2i\phi + R^4. \tag{S5}$$

To obtain the final LG spectrum coefficients for H1, we integrate $q_{\text{H1,H2,H3}}$ over the LG modes at $z = 0$:

$$c_{l,p}(w) = \int\int_{\mathbb{R}^2} d^2\boldsymbol{R}\, \text{LG}_{-l,p}(R, \varphi, z = 0, w) \text{LG}_{0,0}(R, \varphi, z = 0, w) q_{\text{H1,H2,H3}}(R, \varphi, z = 0). \tag{S6}$$

Fig. S1. Optical knot basis. Field phases and corresponding amplitude distributions (shown in the insets) at $z = 0$, along with the LG spectra and 3D visualizations of the singularity line structures for each optical Hopf link (H1–H9) and optical trefoil knots (T1, T2) in Alphabet-11. The names are displayed above each panel. The color bar for all amplitude distributions in the insets is shown in the first panel and remains constant for all subsequent panels.

The parameter w defines the shape of the optical knot and must be selected within a specific range to preserve the desired topological configuration [1,2]. The values for w are: H1: $w = 1.4$, H2: $w = 1.6$, H3: $w = 1.8$.

H4 Hopf Link

To construct H4, we start with the H2 Hopf link ($w = 1.6$) and modify the inverse stereographic projection (S4) by rotating the two loops of the link around the x-axis in opposite directions: one loop is rotated by $\theta = \pi/6$, and the other loop is rotated by $\theta = -\pi/6$. The transformation is given by:

$$x \to x, \qquad y \to y\cos(\theta) - z\sin(\theta), \qquad z \to y\sin(\theta) + z\cos(\theta). \tag{S7}$$

As a result, the final Milnor map in the $z = 0$ plane takes the form:

$$q_{\mathrm{H4}}(R,\varphi,z=0) = (+R^4 + R^2[-2-2e^{-i2\varphi}-2e^{2ia}-2(-e^{-i2\varphi}+e^{i2\varphi})\cos(\theta_1) + 2(-e^{-i2\varphi}+e^{i2\varphi})\sin(\theta_1) - 2\sin(\theta_1) + e^{-i2\varphi}\sin(\theta_1) + e^{i2\varphi}\sin(\theta_1)]). \quad \text{(S8)}$$

For a detailed derivation, see Appendix B in [3].

H5 Hopf Link

To construct the H5 Hopf link, we start with the classical H2 knot with $w = 1.6$ and modify the inverse stereographic projection (S4) in two steps. First, we rotate one of the loops around the z-axis by an angle $\alpha = \pi/6$, applying the transformation:

$$x \to x\cos(\alpha) - y\sin(\alpha), \qquad y \to x\sin(\alpha) + y\cos(\alpha), \qquad z \to z. \quad \text{(S9)}$$

To compensate for the vertical displacement caused by this transformation, one lobe is shifted upward by $z_1 = 0.3$, while the other is shifted downward by $z_2 = -0.3$. The Milnor polynomial at $z = 0$ for this modified Hopf link takes the form:

$$q_{\mathrm{H5}}(R,\varphi,z=0) = \left(1 + R\left(2e^{i\varphi} - 2e^{i(\varphi+\alpha)} - 4ie^{i(\varphi+\alpha)}z_1 + 2e^{i(\varphi+\alpha)}{z_1}^2 + 4ie^{i\varphi}z_2 - 2e^{i\varphi}{z_2}^2\right) + R^2\left(-2 - 4e^{i(2\varphi+\alpha)} - 2iz_1 + {z_1}^2 - 2iz_2 + {z_2}^2\right) + R^3\left(-2e^{i\varphi} + 2e^{i(\varphi+\alpha)}\right) + R^4 + 2iz_1 - {z_1}^2 + 2iz_2 - 4z_1z_2 - 2i{z_1}^2z_2 - {z_2}^2 - 2iz_1{z_2}^2 + {z_1}^2{z_2}^2\right). \quad \text{(S10)}$$

For a detailed derivation, see [3].

H6 Hopf Link

Hopf link H6 is a numerically optimized Hopf link designed to improve stability in turbulent environments. A detailed description of the optimization approach used for this knot is provided in [3].

H7 Hopf Link

Hopf link H7 is constructed by shifting the two loops along the z-axis in opposite directions by 0.3. The mathematical expression for this link is essentially the same as for H5, but without any rotation, meaning the angle α is set to 0.

H8 Hopf Link

Hopf link H8 is obtained by keeping one of the loops unchanged while rotating the other by an angle $\theta = \pi/9$ along the x-axis. The procedure follows the same approach as in H4, but with a different rotation angle applied to only one of the lobes.

H9 Hopf Link

Hopf link H9 is constructed using the same method as H5, with the only difference being the rotation angle α. For H9, the angle is $\alpha = \pi/12$ instead of $\pi/6$ as used in H5. The shifts along the z direction remain the same at ± 0.3.

T1 Trefoil Knot

Trefoil T1 is a classical isolated optical knot obtained without modifications to the stereographic projection, similar to Hopf links H1, H2, and H3. The Milnor map for this trefoil is defined as:

$$q_{\mathrm{T1}} = (u - v^{1.5})(u + v^{1.5}), \quad \text{(S11)}$$

which, upon eliminating the shared denominator, transforms at $z = 0$ into the following expression:

$$q_{\mathrm{T1}}(R,\varphi,z=0) = 1 - R^2 - 8R^3\exp 3i\phi - R^4 + R^6. \quad \text{(S12)}$$

Parameter w is set to $w = 1.2$.

T2 Trefoil Knot

Trefoil T2 is a numerically optimized knot designed for improved stability in a turbulent environment. The optimization approach follows the same method used for H6 and is described in detail in [3].

Table S1. Coefficients $c_{l,p}$ for Alphabet-11 Knots.		
Knot	Indexes (l, p) in $c_{l,p}$ coefficients	Corresponding coefficients $c_{l,p}$
H1	(0, 0), (0, 1), (0, 2), (2, 0)	2.60, -6.25, 4.19, -6.05
H2	(0, 0), (0, 1), (0, 2), (2, 0)	2.98, -7.00, 4.35, -4.81
H3	(0, 0), (0, 1), (0, 2), (2, 0)	3.25, -7.44, 4.40, -3.84
H4	(-2, 0), (0, 0), (0, 1), (0, 2), (2, 0)	-0.64, 2.00, -7.32, 5.76, -2.97
H5	(0, 0), (0, 1), (0, 2), (1, 0), (1, 1), (2, 0)	2.89, -6.53, 4.01, 0.60, -0.35, -3.91
H6	(0, 0), (0, 1), (0, 2), (2, 0)	2.96, -6.22, 4.74, -5.48
H7	(0, 0), (0, 1), (0, 2), (1, 0), (2, 0)	3.09, -7.00, 4.29, -1.22i, -4.65
H8	(-1, 0), (-1, 1), (0, 0), (0, 1), (0, 2), (1, 0), (1, 1), (2, 0)	-0.97, 0.86, 2.55, -6.99, 4.82, 0.97, -0.86, -4.25
H9	(0, 0), (0, 1), (0, 2), (1, 0), (1, 1), (2, 0)	3.02-0.15i, -6.82+0.35i, 4.18-0.21i, 0.18-2.27i, -0.08+0.98i, -4.43-0.95i
T1	(0, 0), (0, 1), (0, 2), (0, 3), (3, 0)	1.71, -5.66, 6.38, -2.3, -4.36
T2	(0, 0), (0, 1), (0, 2), (0, 3), (3, 0)	1.29, -3.95, 7.49, -3.28, -3.98

S2. Examples of Elements of Alphabet-11 Optical Knots in Turbulence of Different Strengths

Figure S2 presents three different Hopf links – H4, H6, and H8 – under three turbulence strengths studied in this paper: $\sigma_R^2 = 0.05$, $\sigma_R^2 = 0.15$, and $\sigma_R^2 = 0.25$. By comparing the field distributions, LG spectra, and 3D structures of these knots after turbulence with their unperturbed counterparts shown in Fig. S1, we can observe how turbulence affects the propagation of optical knots.

In Fig. S2a, representing the weakest turbulence case, the overall 3D structures of the knots remain nearly unchanged, though spectral broadening is already noticeable. As turbulence strength increases, Fig. S2b shows more pronounced spectral crosstalk, and the shape of the knots becomes increasingly distorted, though their initial structures are still recognizable. Under the strongest turbulence, as shown in Fig. S2c, both spectral broadening and structural deformation become more significant, making the original knot configurations more challenging to identify.

S3. Full Spectrum-Based MSE Classification of Alphabet-11.

Figure S3 presents a confusion matrix and the corresponding average classification accuracy calculated over three turbulence strengths ($\sigma_R^2 = 0.05, 0.15, 0.25$) for Alphabet-11 knots, using the complete spectral range for each knot realization – specifically, the entire LG spectrum with indices $l = -10$ to 10 and $p = 0$ to 10. The total classification accuracy in this case is 28.2%, which is notably lower than the accuracy obtained with the alternative MSE-based approach discussed in the main text (34.8%), where only the basis LG modes were used for knot classification. Notably, in both classification methods, the turbulence-optimized optical Hopf link H5 exhibits the highest classification accuracy and most consistent recognition across the entire Alphabet-11 basis.

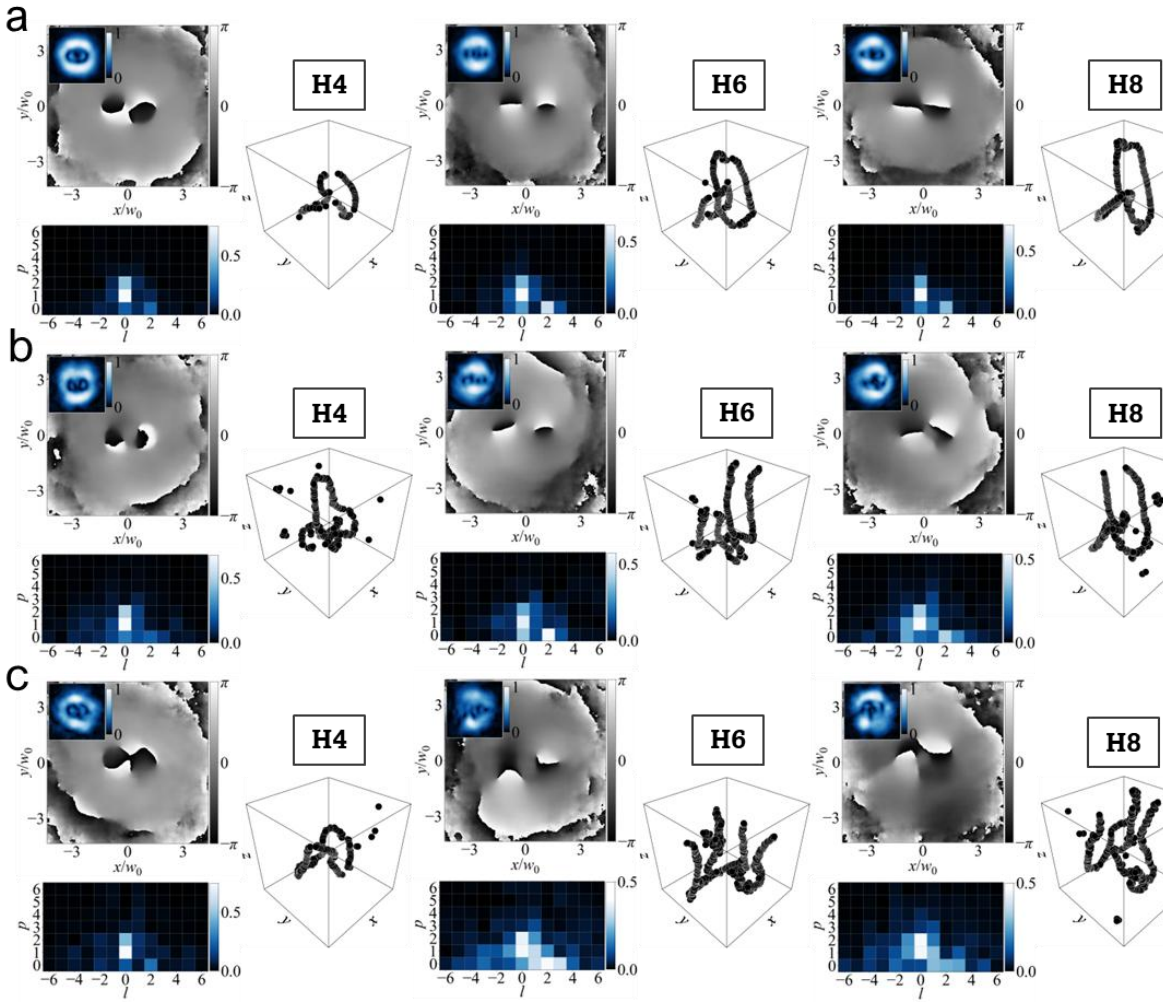

Fig. S2. Demonstration of the effect of turbulence of different strengths on optical knots based on the examples of H4, H6, and H8. Phase distributions and corresponding fields amplitudes (insets) at $z = 0$, along with the LG spectrum and 3D singularity line structures for three optical Hopf links: H4 (first column), H6 (second column), and H8 (third column), after exposure to turbulence. (a) Knot structures after turbulence with $\sigma_R^2 = 0.05$. (b) Knot structures after turbulence with $\sigma_R^2 = 0.15$. (c) Knot structures after turbulence with $\sigma_R^2 = 0.25$.

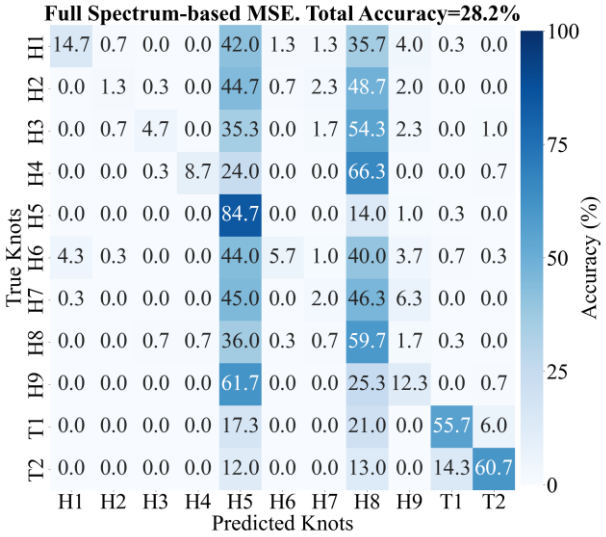

Fig. S3. Confusion matrix for the classification of the alphabet-11 knot basis after turbulence using spectrum-based MSE classification with all the modes in spectra. The confusion matrix displays the percentage of each optical knot being classified into different knot types after experiencing turbulence. The confusion matrix incorporates data from all three turbulence strengths $\sigma_R^2 = 0.05, 0.15, 0.25$.

S4. MSE Classification of Alphabet-11 at Different Turbulence Strengths

Figure S4 presents confusion matrices for three different turbulence strengths using two MSE-based classification methods: spectrum-based (panel a) and shape-based MSE (panel b). The first column corresponds to the weakest turbulence with $\sigma_R^2 = 0.05$, the second column represents moderate turbulence with $\sigma_R^2 = 0.15$, and the third column shows results for the strongest turbulence with $\sigma_R^2 = 0.25$.

For all three turbulence levels, the shape-based MSE method consistently achieves higher classification accuracy compared to the spectrum-based approach. In the weakest turbulence case, the shape-based method reaches an accuracy of 72.1%, whereas the spectrum-based approach achieves 50.3%, highlighting the advantage of incorporating spatial structure information in turbulent environments.

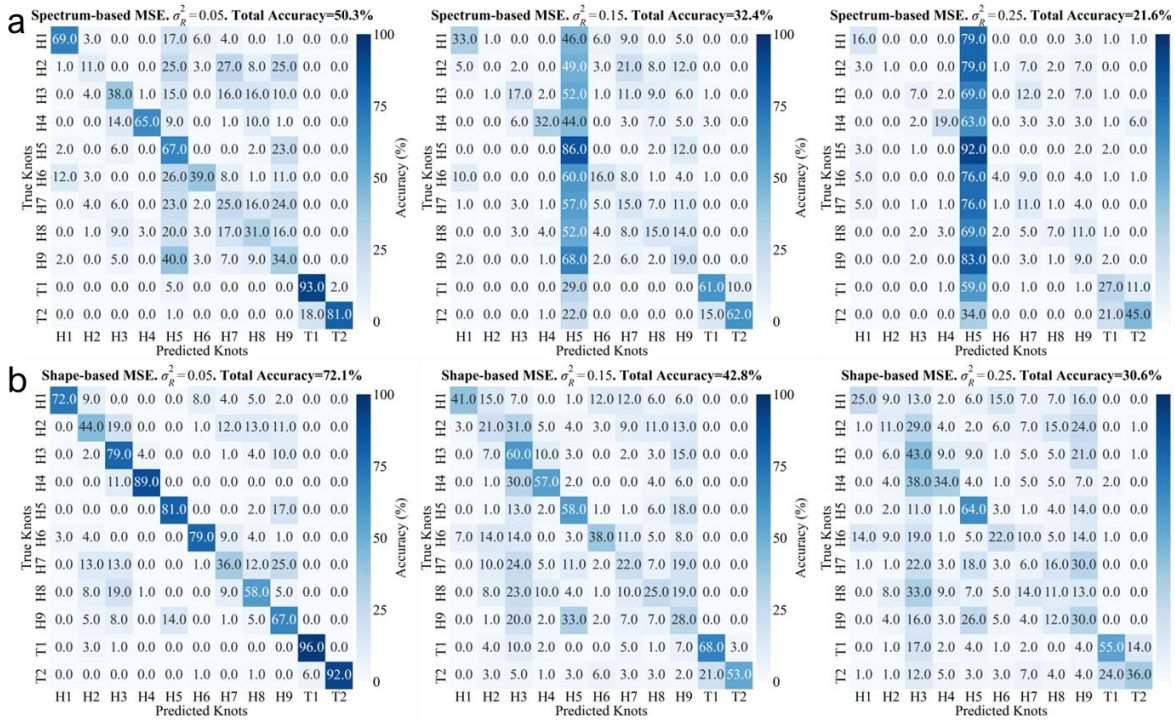

Fig. S4. Confusion matrices for the MSE-based classification of the optical knot basis after turbulence. The confusion matrices display the percentage of each optical knot being classified into different knot types after experiencing turbulence. (a) Spectrum-based MSE classification. (b) Shape-based MSE classification. The average classification accuracy is indicated at the top of each panel. The first column shows the weakest turbulence case $\sigma_R^2 = 0.05$, the second column is the moderate case σ_R^2=0.15, while the last column is the strongest turbulence case $\sigma_R^2 = 0.25$.

S5. Spectrum-Based Machine Learning Algorithms

Models' architecture

We developed two spectrum-based machine learning models for classifying structured optical fields: a fully connected neural network (FCN) and a convolutional neural network (CNN). Both models process spectral data but differ in how they handle spatial correlations. The FCN treats spectral coefficients as a feature vector, whereas the CNN interprets the spectrum as a two-dimensional representation, capturing relationships between spectral components through convolutional layers. Below, we describe the architecture and training procedures for these models.

The Spectrum-Based FCN for Alphabet-11 is a fully connected neural network designed for spectral classification. It consists of an input layer with 91 features, followed by a sequence of fully connected layers, each incorporating batch normalization, ReLU activation, and dropout for regularization. The final architecture includes 12 fully connected layers. Most layers contain 1024 units, while the first and last layers contain 128 units. Each layer is followed by batch normalization, ReLU activation, and a dropout layer with a rate of 25% to prevent overfitting. The output layer applies a linear transformation to map the features to 11 classes. The total number of trainable parameters is 10,795,531. Table S2 provides a detailed description of the FCN model.

Table S2. Overview of the Spectrum-Based FCN Model Architecture

Layer	Description	Output Shape
Input Layer	7x13 weights	[−1,91]
Fully Connected 1	Linear (91 → 128) + BN + ReLU + Dropout(0.25)	[−1,128]

Fully Connected 2	Linear (128 → 1024) + BN + ReLU + Dropout(0.25)	[−1,1024]
Fully Connected 3-11	Linear (1024 → 1024) + BN + ReLU + Dropout(0.25)	[−1,1024]
Fully Connected 12	Linear (1024 → 128) + BN + ReLU + Dropout(0.25)	[−1,128]

The Spectrum-Based CNN for Alphabet-11 is a convolutional neural network designed for spectral classification. It processes two-dimensional spectral data using convolutional layers to capture spatial dependencies. The architecture consists of two convolutional stages, each containing three convolutional layers with kernel size $k = 3$, stride $s = 1$, and padding $p = 1$. Each layer is followed by batch normalization and ReLU activation functions. A 2D max-pooling layer ($k = 2$, $s = 2$, $p = 1$) is applied after each stage to reduce the spatial dimensions. The extracted features are flattened and passed through two fully connected layers: one with 256 units (ReLU activation) and a final layer with 11 output units. The total number of trainable parameters is 311,435. See Table S3 for details.

Table S3. Overview of the Spectrum-Based CNN Model Architecture

Block	Layers & Parameters	Output Shape
Block 1	3 × [Conv2D($k = 3,\ s = 1,\ p = 1$)) + BN + ReLU] + MaxPool2D($k = 2, s = 2, p = 1$)	[−1,32,4,7]
Block 2	3 × [Conv2D($k = 3,\ s = 1,\ p = 1$)) + BN + ReLU] + MaxPool2D($k = 2, s = 2, p = 1$)	[−1,64,6,6,6]
Fully Connected	Flatten → Linear(64×3×4→256) + ReLU → Linear(256 →11)	[−1,11]

Training hyperparameters

To optimize model performance, a comprehensive hyperparameter grid search was conducted across a range of values, including learning rate, decay schedule, dropout rates, layer sizes, and kernel configurations. The explored parameter ranges are summarized in Table S4.

Table S4. Hyperparameter Search Space for Spectrum-Based FCN and CNN

Hyperparameter	Value
Learning rates	10^{-3}, 5×10^{-3}, 10^{-4}, 5×10^{-4}, 10^{-5}, 5×10^{-5}, 10^{-6}
Decay epochs	0,5,10,25,50
Decay factors	0.1,0.2,0.5
Batch sizes	32,64,128
Dropout Rates (FCN)	0%, 10%, 25%, 50%
Number of Layers (FCN)	5, 7, 10, 12
Layer Sizes (FCN)	256, 512, 1024, 1280
Kernel Sizes (CNN)	3, 5
Number of Stages (CNN)	1, 2, 3
Stage Lengths (CNN)	2, 3
Number of Epochs	25, 50, 75, 100

Following the evaluation, the final selected hyperparameters provided optimal performance for both models. These values are presented in Table S5. These optimized hyperparameters ensured stable training and high classification accuracy across all turbulence strengths.

Table S5. Final Training Hyperparameters for Spectrum-Based FCN and CNN		
Hyperparameter	FCN Value	CNN Value
Loss Function	Cross-Entropy	Cross-Entropy
Optimizer	Adam	Adam
Initial LR	5×10^{-4}	2×10^{-5}
LR Decay Epochs	25	25
LR Decay Factor	0.2	0.2
Batch Size	64	64
Dropout Rate	25%	-
Number of Epochs	50	50

S6. Machine Learning Classification of Alphabet-11 at Different Turbulence Strengths

Figure S5 presents confusion matrices for three different turbulence strengths using two machine learning classification methods: (a) spectrum-based FCNN and (b) shape-based CNN. The first column corresponds to the weakest turbulence ($\sigma_R^2 = 0.05$), the second column represents moderate turbulence ($\sigma_R^2 = 0.15$), and the third column shows results for the strongest turbulence ($\sigma_R^2 = 0.25$).

Across all turbulence levels, the shape-based CNN consistently outperforms the spectrum-based approach in classification accuracy. In the weakest turbulence case, the shape-based method achieves 97.7% accuracy, compared to 79.7% for the spectrum-based approach, demonstrating the benefits of incorporating spatial structure information. Even under the strongest turbulence conditions, the shape-based method maintains an accuracy of 75.4%, while the spectrum-based approach drops to 53.8%, further emphasizing the robustness of shape-based features in turbulent environments.

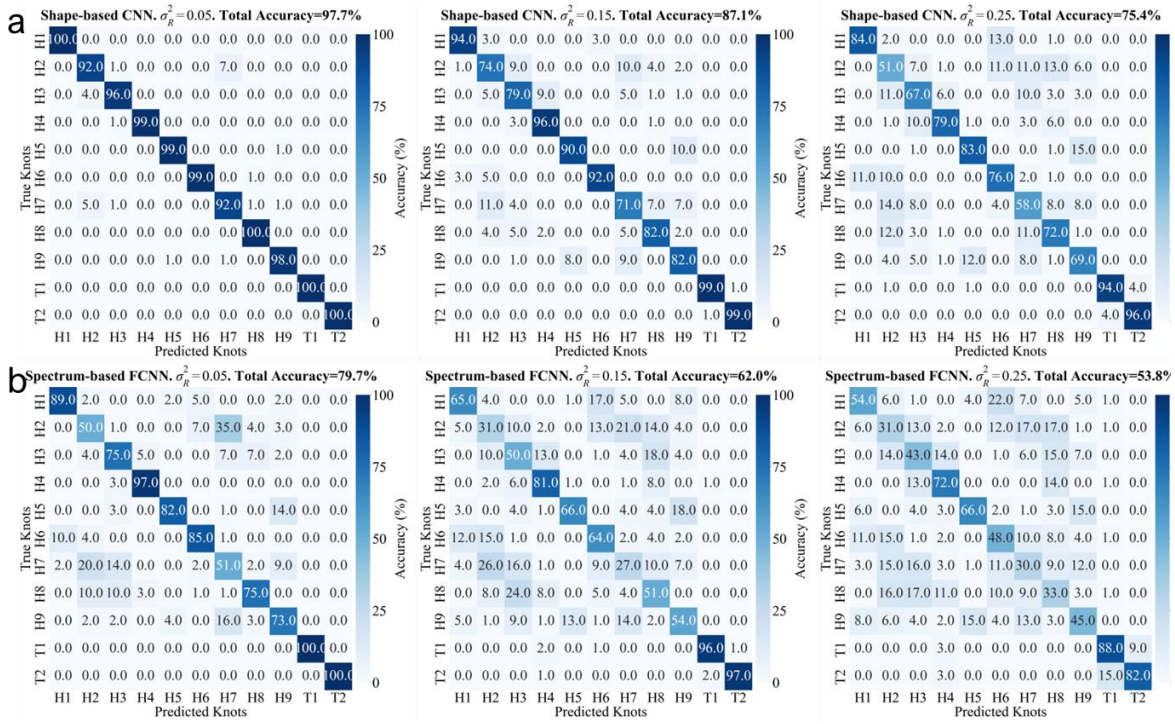

Fig. S5. Confusion matrices for machine learning classification of the optical knot basis after turbulence. The confusion matrices illustrate the classification performance of Alphabet-11 knots under three turbulence strengths using two machine learning models: (a) Spectrum-based FCNN and (b) Shape-based CNN machine learning classification. The first column corresponds to the weakest turbulence case ($\sigma_R^2 = 0.05$), the second column represents moderate turbulence ($\sigma_R^2 = 0.15$), and the last column shows results for the strongest turbulence ($\sigma_R^2 = 0.25$). The average classification accuracy is indicated at the top of each panel.

S7. Shape-Based Machine Learning Algorithms

Models' architecture

We developed two 3D convolutional neural network (CNN) models to classify structured optical fields: one for Alphabet-11 (optical knots) and another for Alphabet-81 (Flower beams). Both models share the same architecture, differing only in the number of output classes, which affects the final fully connected layer. Below, we describe the models' architecture along with the training and tuning procedures.

The Shape-Based CNN for Alphabet-11 consists of two main processing blocks. The first block comprises three consecutive 3D convolutional layers, each with a kernel size of 3, a stride of 1, and a padding of 1, followed by batch normalization (BN) and a ReLU activation function to ensure stable and efficient training. A 3D max-pooling layer with a kernel size of 2, a stride of 2, and a padding of 1 reduces spatial dimensions. The second block mirrors this structure but increases the kernel size to 5, capturing more complex spatial features. The extracted feature maps are flattened and passed through two fully connected layers: one with 256 neurons and ReLU activation, followed by a final classification layer with 11 output neurons, corresponding to the number of optical knot classes. The model has 4,879,051 trainable parameters and an input volume size of 32 × 32 × 32. Table S6 provides a detailed description of the architecture.

The Shape-Based CNN for Alphabet-81 follows the same architecture but adjusts the final fully connected layer to have 81 output neurons, aligning with the classification task for Flower beams. This modification increases the total number of trainable parameters to 4,897,041.

Table S6. Shape-based CNN for Alphabet-11

Block	Layers & Parameters	Output Shape

Convolutional Block 1	3 × [Conv3D($k = 3, s = 1, p = 1$)) + BN + ReLU] + MaxPool3D($k = 2, s = 2, p = 1$)	[−1,32,17,17,17]
Convolutional Block 2	3 × [Conv3D($k = 5, s = 1, p = 1$)) + BN + ReLU] + MaxPool3D($k = 2, s = 2, p = 1$)	[−1,64,6,6,6]
Fully Connected Layer	Flatten → Linear(64×6×6×6→256) + ReLU → Linear(256 →11)	[−1,11]

Training hyperparameters

To determine the most effective training configuration, a grid search was conducted across multiple hyperparameters. The parameters explored in this search are summarized in Table S7.

Table S7. Hyperparameter Search Space	
Hyperparameter	Value
Learning rates	10^{-3}, 5×10^{-3}, 10^{-4}, 5×10^{-4}, 10^{-5}, 5×10^{-5}, 10^{-6}
Decay epochs	0,5,10,25,50
Decay factors	0.1,0.2,0.5
Batch sizes	32,64,128
Number of epochs	25, 50, 75, 100

After evaluating these configurations, the optimal hyperparameters were selected based on the best performance achieved on Alphabet-11 and subsequently applied to Alphabet-81. These settings provided a balance between training stability and performance, ensuring consistent classification accuracy across different structured optical fields. The final hyperparameters are listed in Table S8. By using these optimized parameters, both Shape-Based CNN models effectively classify structured optical fields while maintaining robust generalization across different turbulence conditions.

Table S8. Final Training Hyperparameters	
Hyperparameter	Value
Loss Function	Cross-Entropy
Optimizer	Adam
Initial LR	10^{-5}
LR Decay Epochs	25
LR Decay Factor	0.2
Batch Size	64

S8. Full Spectrum-Based MSE Classification of Alphabet-11

This section presents the results of the classification using the spectrum-based convolutional neural network (CNN). The confusion matrix, combining all turbulence strengths, is shown in Fig. S6. The total classification accuracy achieved with this method is 62.6%, which is comparable to the accuracy of the spectrum-based fully connected neural network (FCNN) used in the main text (65.2%).

The slight difference in accuracy may be attributed to classification errors and potential sub-optimal training procedures for the spectrum-based MSE method, where the best global parameters might not have been fully optimized

during grid search. Overall, the similarity in accuracy between these methods suggests that we are approaching the optimal classification performance achievable using spectral data.

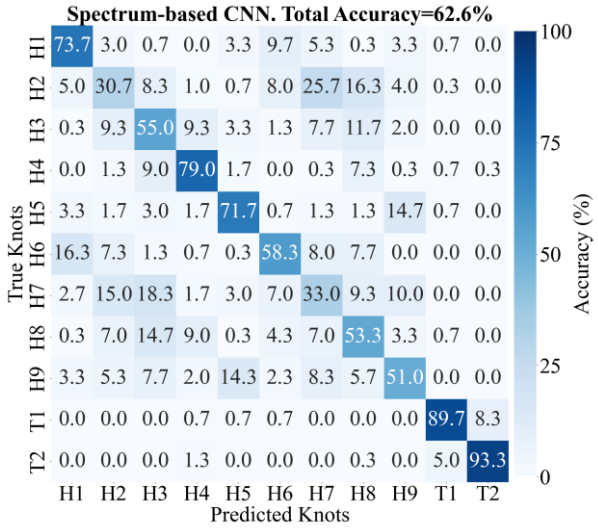

Fig. S6. Confusion matrix for the classification of Alphabet-11 knots after turbulence across all turbulence strengths using the spectrum-based CNN. The confusion matrix displays the percentage of each optical knot classified into different knot types after experiencing turbulence, using the spectrum-based CNN machine learning algorithm. The average classification accuracy is 62.6%. The matrix incorporates data from all three turbulence strengths: $\sigma_R^2 = 0.05$, $\sigma_R^2 = 0.15$, and $\sigma_R^2 = 0.25$.

S9. Generation of Flower Beams

To determine the coefficients $c_{l,p}$ of the LG spectrum of Flower beams, analogous to equation (S1), we begin with the braid representation, illustrated in Fig. S7a. In this approach, a periodic complex scalar field with braided zero-lines is embedded in a cylinder of height of 2π, within the coordinate system (x', y', h) defined as:

$$\begin{cases} x'(h) = \ cos(sh), \\ y'(h) = \ sin(sh), \end{cases} \tag{S13}$$

where the parameter s determines the number of petals in the Flower beam. Introducing the substitution $u(x', y') = x' + iy'$, and $v(h) = e^{ih}$, we obtain the Milnor map expression for a Flower beam with $s = 4$ full-sized petals:

$$q_{\text{flower}}(u, v) = u - v^s. \tag{S14}$$

By utilizing the inverse stereographic projection (S4) and multiplying the expression by $(R^2 + z^2 + 1)^s$ at the $z = 0$ plane, we get the Milnor polynomial at the $z = 0$ plane:

$$Q_{\text{flower}}(R, \varphi, z = 0) = -(2e^{if}R)^s + (-1 + R^2)(1 + R^2)^{-1+s}. \tag{S15}$$

This expression allows us to compute the LG spectrum $c_{l,p}$ using equation (S6). The specific case for $s = 4$ is given by:

$$Q_{\text{flower}}^{s=4}(R, \varphi, z = 0) = -16e^{4if}R^4 + (-1 + R^2)(1 + R^2)^3. \tag{S16}$$

Figure S7b illustrates the corresponding field, its LG spectrum, and the 3D structure. The different colors in the figure represent distinct lobes, allowing us to trace their origins in the braid representation. The numbering of lobes starts at $\varphi = 0$ and progresses counterclockwise. The angular intervals for each petal are defined as follows:

$$\varphi = \begin{cases} \left(-\frac{\pi}{4}, \frac{\pi}{4}\right] : \text{ lobe 1}, \\ \left(-\frac{3\pi}{4}, -\frac{\pi}{4}\right] : \text{ lobe 2}, \\ \left(\frac{3\pi}{4}, \pi\right] \bigcup \left[-\pi, -\frac{3\pi}{4}\right] : \text{ lobe 3}, \\ \left(\frac{\pi}{4}, \frac{3\pi}{4}\right] : \text{ lobe 4}. \end{cases} \tag{S17}$$

The value 2 in Fig. S7b indicates that all petals are of full size. To control petal sizes, we apply the same shape manipulation approach used for optical knots in Alphabet-11, as detailed in [3]. For each lobe defined in equation (S17), its size can be independently adjusted by scaling the stereographic projection (S4) along the x and y axes within the corresponding φ regions. This is achieved by replacing x and y with γx and γy, where γ is the scaling coefficient.

In our study, we consider three scaling cases, denoted in the paper as 2, 1, and 0, with corresponding values of $\gamma = 1, 1.5$, and ∞. In practice, removing a lobe does not require $\gamma = \infty$. With a 256×256 resolution used for field propagation in turbulence, any $\gamma > 50$ produces no noticeable difference, as the lobe effectively disappears. Since γ scales the lobe size $\sim 1/\gamma^2$ in XY plane, decreasing the lobe size infinitely or by a sufficiently large finite factor has the same effect in a limited-resolution setting.

It is worth noting that although we refer to petals of size 1 as "half-sized" petals, their actual size reduction is not exactly by a factor of 2. The scaling in the XY plane follows $\gamma^2 = 2.25$, and while scaling along the z-axis is not explicitly defined in our case, it is still present. The actual petal size along the z-axis depends on its projection from $z = 0$ according to equation (S4) and is indirectly affected by the scaling in the XY plane. The final size of each half-sized petal slightly varies for different Flower beams. For more details, see [3].

An example of a Flower beam 2220, where the first three lobes have full size 2 and the fourth lobe is removed 0, is shown in Fig. S7c. Another example, Flower beam 2120, where the second lobe is scaled by $\gamma = 1.5$ and the fourth lobe is removed 0, is presented in Fig. S7d.

The total number of Flower beam variations is determined by the size variations per petal, raised to the power s. In our case, with three possible sizes (0, 1, and 2) and $s = 4$, the total number of Flower beams is: $3^4 = 81$.

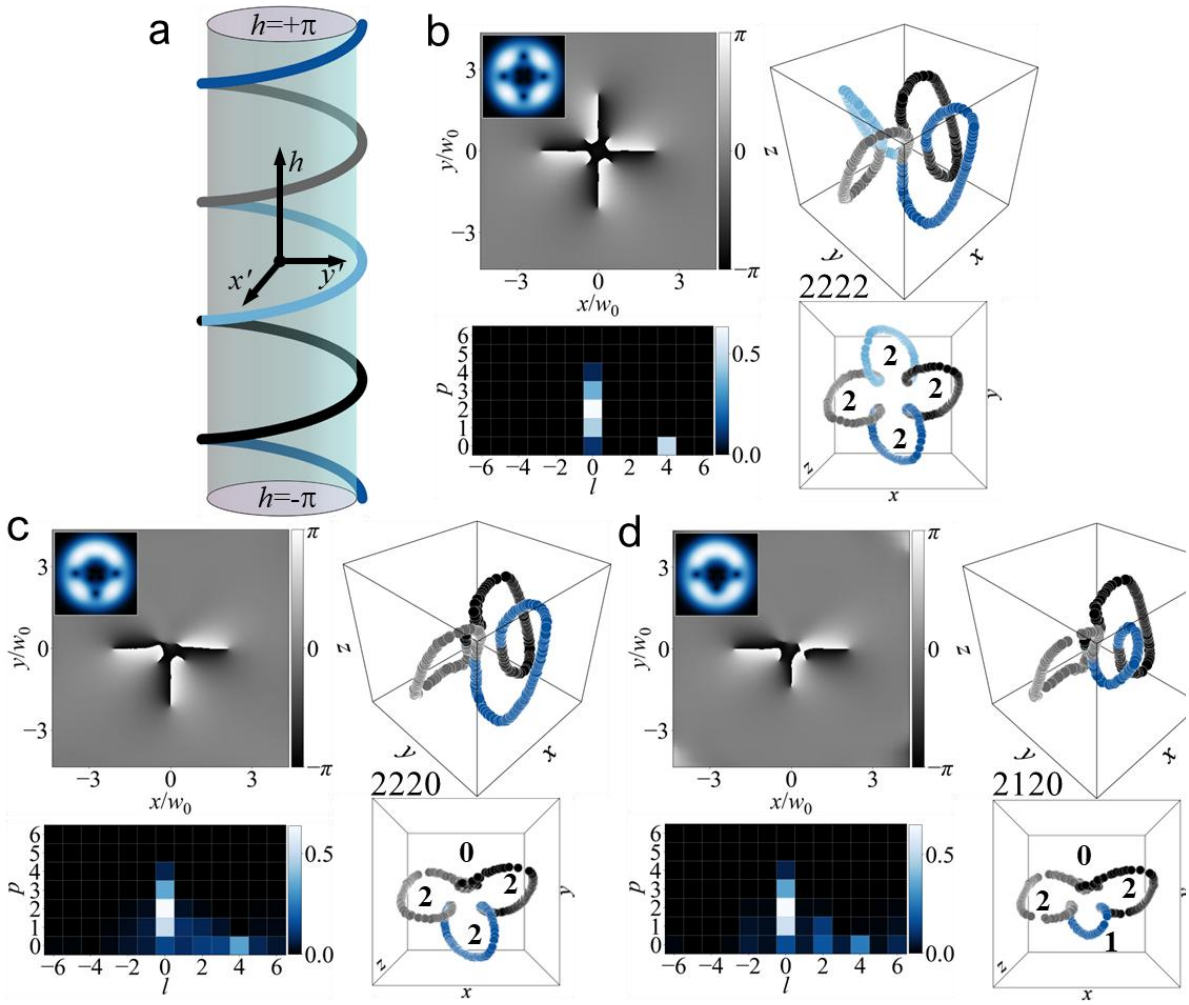

Fig. S7. The concept of creating and modifying Flower beams. (a) Braid representation of an optical Flower beam before applying the stereographic projection. (b) Optical Flower 2222, where all lobes are full size (2 – full size, 1 – half size, 0 – removed). The panels show the amplitude of the field at $z = 0$, with the inset displaying the corresponding phase distribution, the LG spectrum required to generate this structure, and the 3D singularity line shape in two orientations: a general view (top) and a top-down view (bottom right). (c) Optical Flower 2220, where one lobe is removed. (d) Optical Flower 2120, where one lobe is half the size and another is removed. Panels (b), (c), and (d) share the same set of visualizations (phase, amplitude, spectrum, and 3D singularity structures). The lobe colors remain consistent with (a) across all plots.

S10. Additional Examples of Flower Beams with 5 and 6 Petals

The general expression for the Milnor polynomial of a Flower beam with s full-sized petals is given by equation (S15). In the specific cases of $s = 5$ and $s = 6$, the expressions take the following forms:

$$Q^{s=5}_{\text{flower}}(R, \varphi, z = 0) = -32e^{5if}R^5 + (-1 + R^2)(1 + R^2)^4,$$
$$Q^{s=6}_{\text{flower}}(R, \varphi, z = 0) = -64e^{6if}R^6 + (-1 + R^2)(1 + R^2)^5. \quad \text{(S18)}$$

Examples of such Flower beams, including both full-lobed and modified-lobed structures, are illustrated in Fig. S8. As the number of petals increases from 4 to 5 and 6, no significant distortions arise in the Flower beam generation approach (panels b and d). The spectra of these beams broaden, reaching up to $l = 7$ for $s = 5$ and $l = 8$ for $s = 6$, reflecting the increased complexity of the field structure.

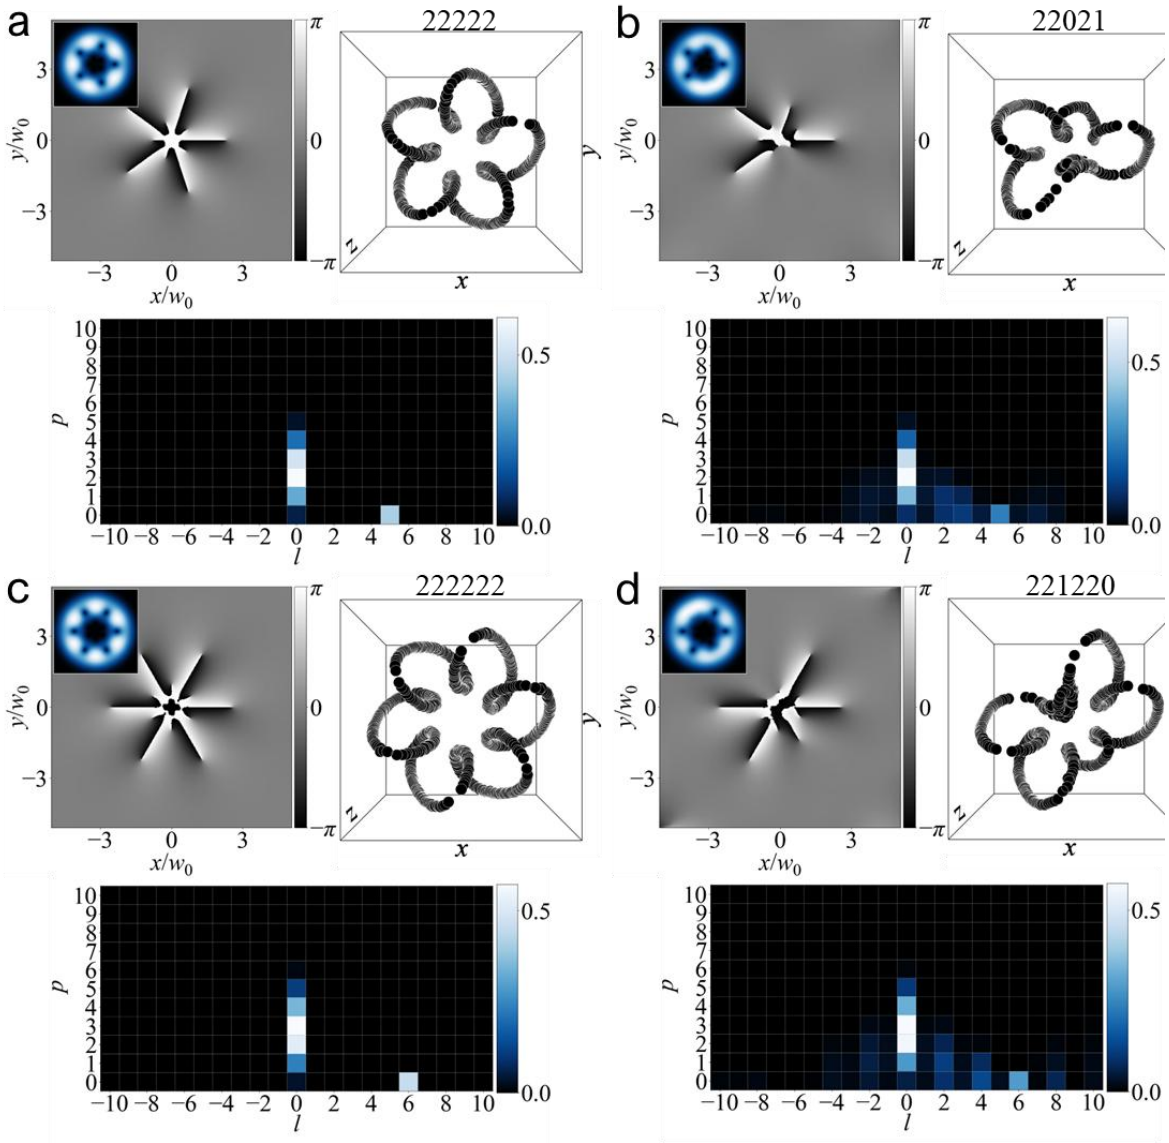

Fig. S8. Examples of Flower beams with five and six petals, including modified petal configurations. This figure presents different Flower beam structures with $s = 5$ (panels a, b) and $s = 6$ (panels c, d). Each panel consists of: (left) the field phase at $z = 0$ with the inset showing the corresponding amplitude distribution, (right) the 3D singularity structure, and (bottom) the LG spectrum. (a) Flower beam with five full-sized petals (22222 configuration). (b) Flower beam with five petals, where the fifth petal is half-sized and the third is removed (22021). (c) Flower beam with six full-sized petals (222222). (d) Flower beam with six petals, where the third and sixth petals are modified (221220).

S11. Flower Beams in Turbulence

Figure S9 presents examples of Flower beams with four petals in two different configurations: 2222 (first column) and 2120 (second column), subject to three different turbulence strengths: $\sigma_R^2 = 0.05, 0.15, 0.25$ (panels a, b, c). For the weakest turbulence case ($\sigma_R^2 = 0.05$), spectral broadening is observed, along with minor deformations in the singularity line structure. As the turbulence strength increases, mode cross-talk in the LG spectrum becomes significantly more pronounced, leading to greater spectral mixing. The singularity structures also experience stronger distortions; however, their overall shape remains recognizable, retaining key features of the unperturbed case shown in Fig. S7. This demonstrates that, despite the effects of turbulence, the fundamental structure of the Flower beams persists, highlighting their robustness in turbulent environments.

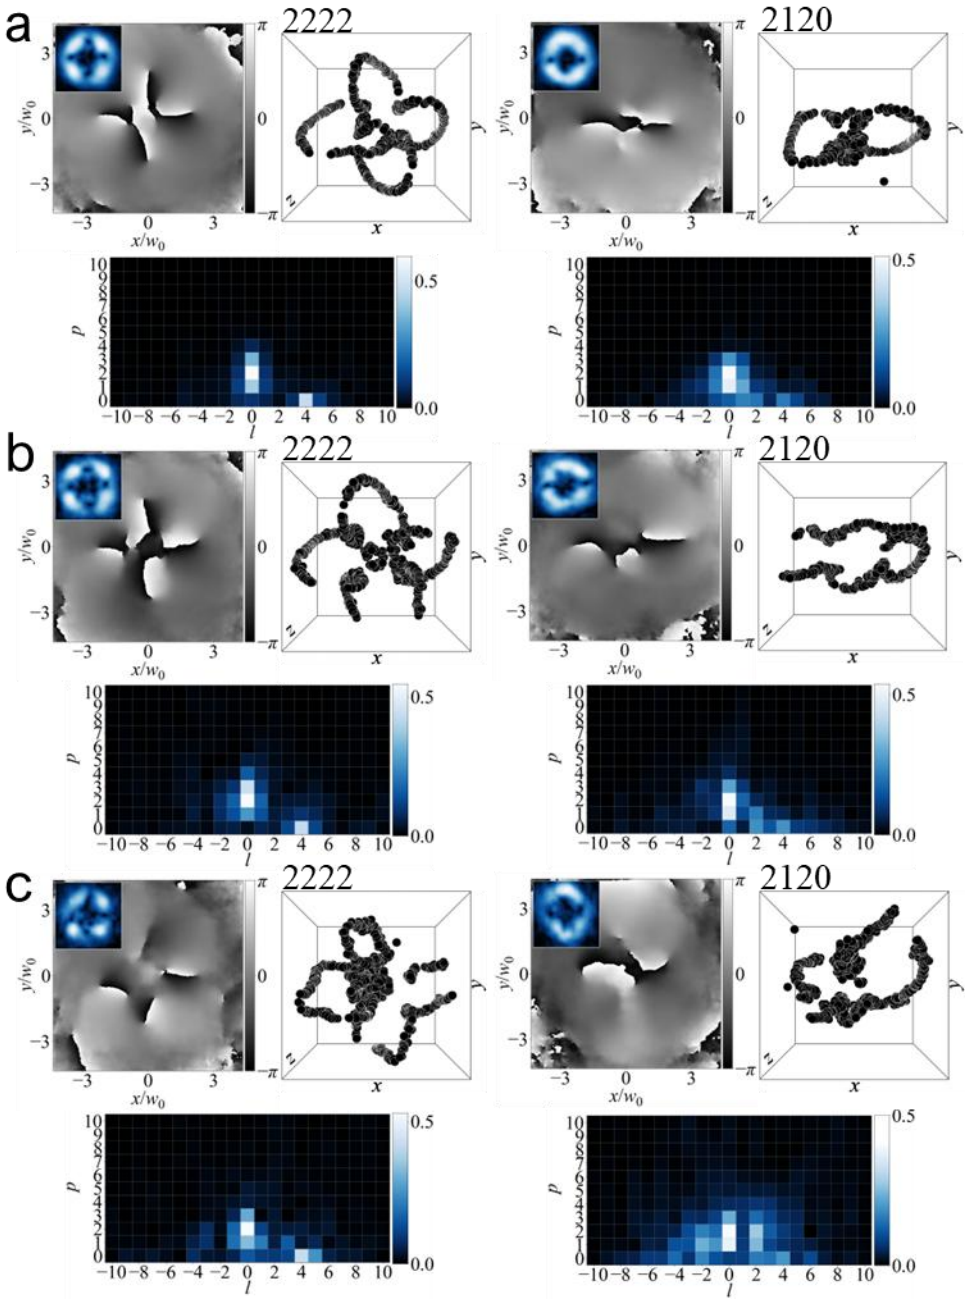

Fig. S9. Flower beams with four petals under different turbulence strengths. This figure illustrates the evolution of 2222 (first column) and 2120 (second column) Flower beams as turbulence strength increases. Each panel includes: (Left) The field phase at $z = 0$, with the inset showing the corresponding amplitude distribution. (Right) The 3D singularity structure of the optical beam. (Bottom) The LG spectrum of the beam after turbulence. (a) Weak turbulence case ($\sigma_R^2 = 0.05$). (b) Moderate turbulence case ($\sigma_R^2 = 0.15$). (c) The strongest turbulence case ($\sigma_R^2 = 0.25$).

S12. Confusion Matrices for Shape-Based CNN Classification on Simulated Data of Alphabet-81 at Different Turbulence Strengths

Figures S10, S11, and S12 present confusion matrices of the optimized shape-based CNN applied to simulated data from Alphabet-81, evaluated under three different turbulence strengths: $\sigma_R^2 = 0.05$, 0.15, and 0.25, respectively. Corresponding total accuracies are 92.7%, 73.6%, and 59.0%.

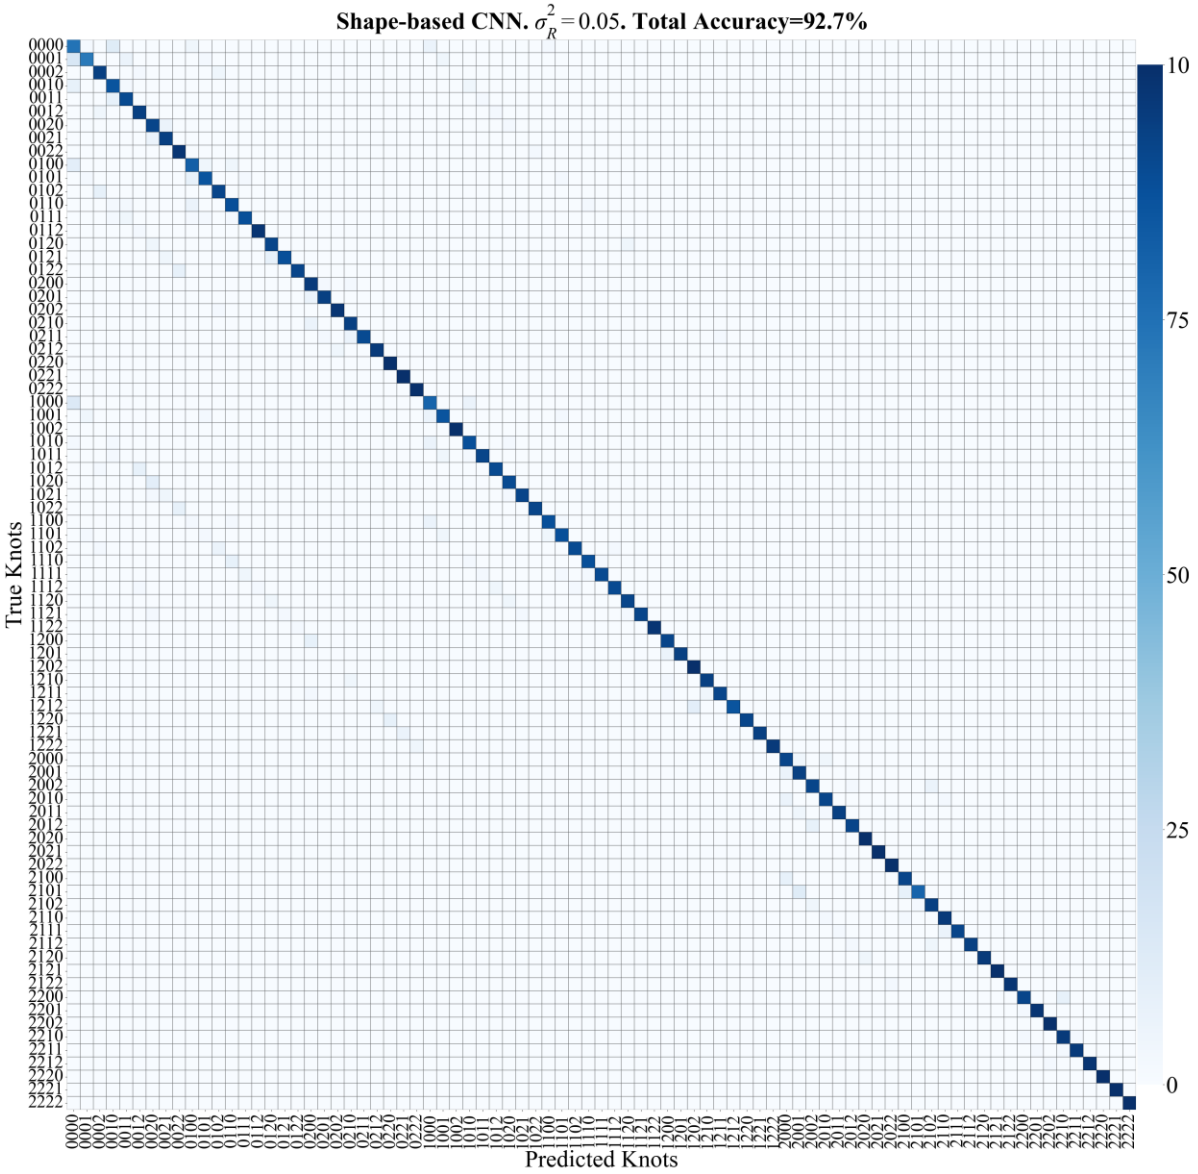

Fig. S10. Confusion matrix for the classification of Flower beams under weak turbulence ($\sigma_R^2 = 0.05$).

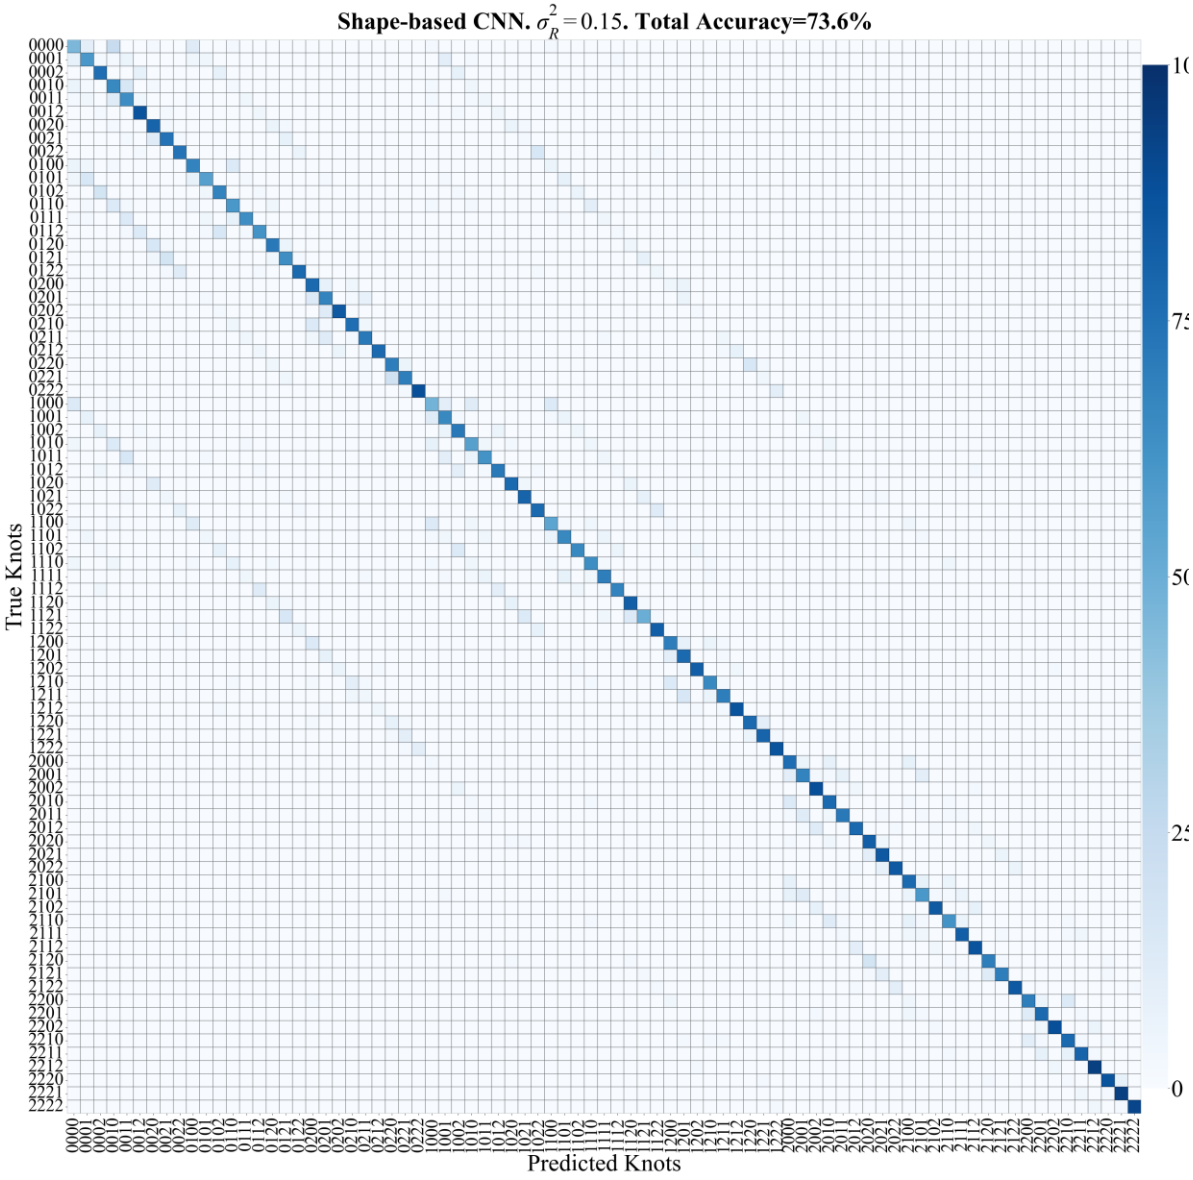

Fig. S11. Confusion matrix for the classification of Flower beams under moderate turbulence ($\sigma_R^2 = 0.15$).

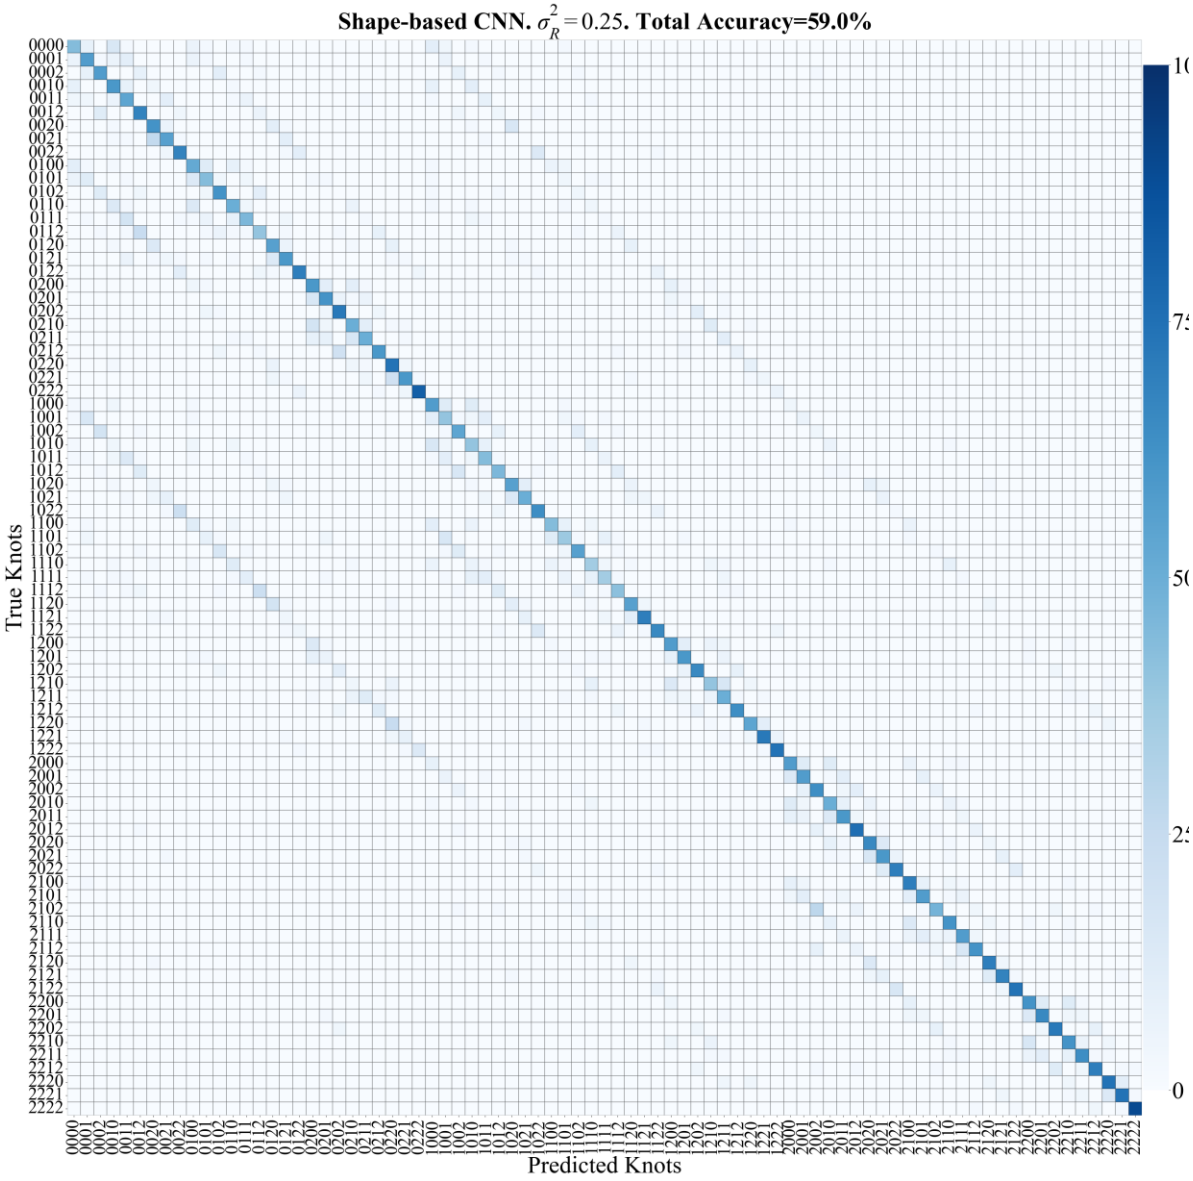

Fig. S12. Confusion matrix for the classification of Flower beams under the strongest turbulence ($\sigma_R^2 = 0.25$).

S13. Estimation of the Turbulence Parameters in the Experimental Setup

To determine the refractive index structure parameter, C_n^2, for our experimental system, we used statistical measurements of the beam wander of a Gaussian beam. The relationship we employed is:

$$C_n^2 = 0.328\langle r_c^2\rangle D^{1/3} L^{-3}, \tag{S19}$$

where D is the beam diameter, L is the turbulence link length, and $\langle r_c^2\rangle = \langle x^2\rangle + \langle y^2\rangle$ represents the variance of the beam wander displacement on the detector. Figure S13a displays the curves of the experimentally estimated C_n^2 obtained from 2000 realizations, along with the corresponding numerical values. Similarly, Figure S13b shows the experimental estimates for the Fried parameter, r_0, overlaid with the numerical predictions.

Our laboratory-scale experimental setup was designed to emulate a realistic free-space optical link. It uses a beam with a 6 mm waist propagating over an equivalent distance of 270 m, achieved through Fresnel scaling to maintain consistent optical properties across both scenarios. Notably, to reproduce comparable scintillation levels in both simulations and experiments, the hot air chamber had to generate stronger turbulence to compensate for the reduced

propagation distance and smaller beam size. For instance, calculating the Rytov variance σ_R^2 to estimate the order of magnitude of C_n^2 for such shorter propagation distance, here $L_{\text{lab}} = 1.5\text{m}$, gives the same order of magnitude as observed in Fig. S13 (a).

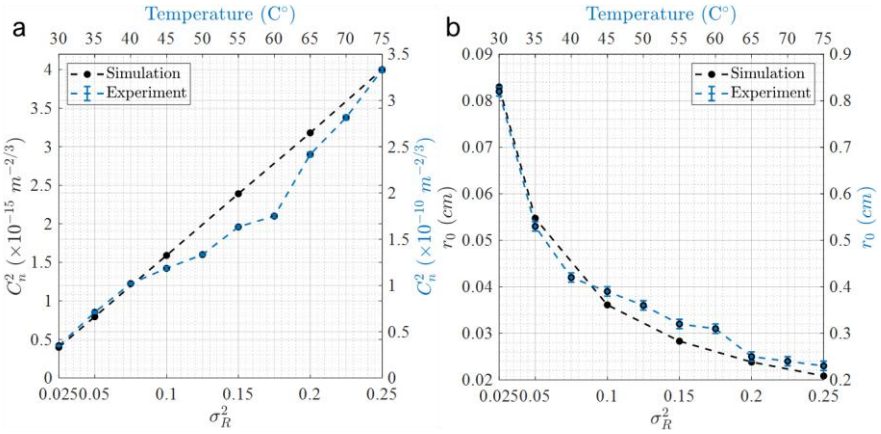

Fig. S13. Comparison of Experimental and Simulated Turbulence Parameters. (a) Experimentally measured refractive index structure parameter C_n^2 (blue) as a function of turbulence chamber temperature, compared with the parameter used in simulations (black), shown against the Rytov variance σ_R^2. A total of 2000 experimental realizations were recorded. (b) Corresponding Fried parameter r_0 values for both simulations (black) and experiments (blue), following the same format as in panel (a).

S14. Fine-Tuning Accuracy of Shape-Based CNN on Alphabet-81 in Experiment

Figure S14 presents the classification accuracy of the shape-based CNN on Alphabet-81 as a function of the number of fine-tuning samples obtained from experimental data. The number of experimental samples used for fine-tuning the model, initially trained exclusively on simulation-generated data, varies from 0 to 10.

The results indicate a steady increase in accuracy as more experimental samples are incorporated into training. However, the rate of improvement slows, eventually approaching a near-saturation level, suggesting that 10 samples are sufficient to achieve accuracy close to its maximum potential. Notably, even with just a single experimental sample, the total accuracy reaches 43% from 17% without extra fine-tuning, demonstrating that the model can generalize well from limited real-world data.

This result can be explained by the fact that, in the experiment, Flower beams without turbulence already exhibit small deviations from their ideal theoretical structures (see Fig. 4 in the main text). As a result, the machine learning algorithm learns these intrinsic deviations inherent to each experimental Flower beam even in the absence of turbulence, refining its classification ability. This adaptation enhances the accuracy of the model trained on simulated data, bringing it closer to optimal performance.

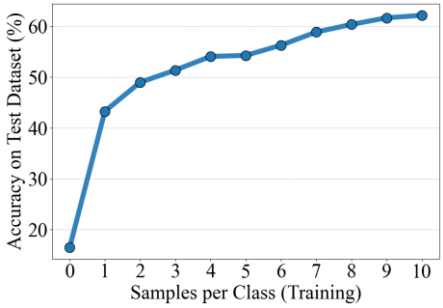

Fig. S14. Accuracy of the shape-based CNN classification on Alphabet-81 in the experiment as a function of the number of fine-tuning samples. This figure illustrates the improvement in classification accuracy achieved by incorporating experimental data into the training process. The x-axis represents the number of fine-tuning samples per class, while the y-axis shows the classification accuracy on the test dataset. As the number of fine-tuning samples increases, accuracy improves but plateaus, indicating diminishing returns with additional samples.

References

1. M. R. Dennis, R. P. King, B. Jack, K. O'Holleran, and M. J. Padgett, "Isolated optical vortex knots," Nat. Phys. **6**, 118–121 (2010).

2. D. G. Pires, D. Tsvetkov, H. Barati Sedeh, N. Chandra, and N. M. Litchinitser, "Stability of optical knots in atmospheric turbulence," Nat. Commun. **16**, 3001 (2025).

3. D. Tsvetkov, D. G. Pires, H. B. Sedeh, and N. M. Litchinitser, "Sculpting isolated optical vortex knots on demand," Photonics Res. **13**, 527–540 (2025).