GSCP-NET: GRAM-STRUCTURED CODED-PROBE RECOVERY OF COMPLEX OPTICAL FIELDS THROUGH ATMOSPHERIC TURBULENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Recovering complex optical fields through atmospheric turbulence is a challenging inverse problem. Our key idea is to cast the complex optical field recovery as parameter estimation, rather than to learn an end-to-end mapping from received to source fields. To make this estimation tractable, we compress the high-dimensional complex field into a low-dimensional latent representation with a variational autoencoder. Inspired by unitary space–time modulation in communications, we exploit the unitary structure of lossless turbulent propagation and propose a latent-variable estimation method that combines multiple unitary coded probes $\mathbf{Q}_k$ with Gram geometry. Under lossless reception, the Gram geometry is independent of the instantaneous turbulence channel, whereas a finite receiver aperture truncates the field. This truncation breaks the ideal invariance, and the joint effect of propagation and truncation reduces to an effective metric $\mathbf{R}$. Since $\mathbf{R}$ can be modeled explicitly, we keep the autoencoder fixed and train only $\mathbf{Q}_k$. Based on this formulation, we propose the Gram-Structured Coded-Probe Network (GSCP-Net) and evaluate it on a 2-km distributed-turbulence channel spanning 42 turbulence–aperture conditions. Compared with end-to-end baselines, GSCP-Net is substantially more accurate under strong turbulence and small apertures, while using orders of magnitude fewer channel-dependent parameters and only a fraction of their training time. Even under $C_n^2 = 3 \times 10^{-13}\,\mathrm{m}^{-2/3}$ with a 140-mm receiver aperture, GSCP-Net reaches a recovery fidelity of 0.660 using 15 probes. Without retraining or channel tracking, the learned probes and fixed receiver keep the fidelity nearly unchanged over 20 coherence times of evolving turbulence. These results establish Gram-structured coded probing as a lightweight, physics-guided approach to robust complex-field recovery across a broad range of turbulence strengths and receiver apertures. Code is available at https://anonymous.4open.science/r/GSCP-Net/.

1 INTRODUCTION

Atmospheric turbulence perturbs the phase of a propagating optical field and, through diffraction, produces coupled amplitude–phase distortions that become increasingly severe as turbulence strengthens. Both model-based and learning-based methods have been developed to mitigate these effects, including direct complex-field reconstruction and compensation of structured optical beams (Liu et al., 2019; Jia et al., 2022; Tan et al., 2022; Xiong et al., 2020). Finite receiver apertures introduce an additional difficulty by truncating the propagated field and altering the geometry of the available measurements (Andrews & Phillips, 2005; Schmidt, 2010).

Many existing learning-based approaches operate with prescribed sensing configurations and do not explicitly exploit the unitary structure of lossless propagation for measurement design. Finite-aperture effects are also less commonly incorporated into the sensing geometry itself. Transmission eigenchannels provide a complementary route for finite apertures (Bachmann et al., 2023), but their construction requires realization-specific characterization of the transmission operator, which is costly in high-dimensional settings.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Motivated by these limitations, we neither learn an end-to-end mapping from received to source fields nor explicitly identify the full transmission matrix. Instead, a variational autoencoder compresses the source field into a low-dimensional latent representation, which is then encoded into multiple unitary coded probes. Under lossless full-field (FF) reception, i.e., when the entire propagated field is received, the common unitary propagation operator preserves the pairwise inner products among the coded fields, and hence their Gram geometry is independent of the instantaneous turbulence realization. Under finite-aperture reception, this exact invariance breaks, and the joint effect of propagation and truncation reduces to a realization-dependent effective metric on the low-dimensional transmit subspace. Our main contributions are as follows:

- We recast turbulence-distorted complex-field recovery as low-dimensional latent-variable estimation using unitary coded probes and Gram geometry (Sec. 3).
- We propose GSCP-Net, which learns only the unitary probe codebook, while its physics-based receiver requires no instantaneous channel state (Sec. 3.4).
- Across 42 turbulence–aperture conditions, GSCP-Net substantially outperforms end-to-end baselines under strong turbulence and small apertures, remains stable under evolving turbulence, and trains in a fraction of their time (Sec. 4).

2 Related Work

Atmospheric-turbulence mitigation includes model-based inversion that exploits wave-propagation physics (Radosevich et al., 2020; Pellizzari et al., 2020) and learning-based approaches that estimate corrective quantities or reconstruct transmitted fields directly. Much of the latter literature has focused on structured optical beams: convolutional neural networks (CNNs) have been used to estimate turbulence distortions or compensate orbital angular momentum (OAM) channels (Liu et al., 2019; Xiong et al., 2020), while diffractive neural networks enable all-optical compensation with low latency (Jia et al., 2022). More recently, Tan et al. (2022) reconstructed both amplitude and phase of general transmitted fields, extending beyond predefined beam families.

To reduce the dimensionality of general complex-field recovery, we use a variational autoencoder (VAE) (Kingma & Welling, 2014) to learn a compact source representation that is subsequently coupled with coded measurement design.

Joint optical encoding and computational reconstruction have also been widely studied in computational imaging (Chang & Wetzstein, 2019; Bacca et al., 2021). Atmospheric propagation, however, introduces a realization-dependent operator between the encoder and receiver. Finite-aperture transmission eigenchannels exploit singular modes of a measured turbulent transmission matrix to enhance transmitted power (Bachmann et al., 2023), but require realization-specific channel characterization and target transmission efficiency rather than source recovery. In contrast, inspired by unitary space–time modulation (Hochwald & Marzetta, 2000), GSCP-Net exploits the unitary structure of lossless diffraction and phase-only turbulence to construct coded measurements whose pairwise Gram geometry is channel-invariant under FF reception. For finite apertures, the same formulation is retained through a condition-dependent mean effective metric, while only the unitary probe codebook is optimized. This provides a common framework for evaluating recovery under both static and dynamically evolving turbulence.

3 Problem Formulation

Throughout this paper, $\mathbb{R}$ and $\mathbb{C}$ denote the fields of real and complex numbers, respectively, and i denotes the imaginary unit. $(\cdot)^T$ and $(\cdot)^H$ denote the transpose and conjugate transpose, $\mathbf{I}_n$ the $n \times n$ identity matrix, and $\mathbf{0}$ a conformable zero matrix or vector. The mappings $\mathcal{E}$ and $\mathcal{D}$ denote the source encoder and decoder, and $\mathbb{E}[\cdot]$ denotes expectation over the relevant distribution. For a matrix $\mathbf{A}$, $\mathrm{rank}(\mathbf{A})$ and $\ker(\mathbf{A})$ denote its rank and null space, while $\dim(\mathcal{S})$ denotes the dimension of a vector space $\mathcal{S}$. The quantities $\sigma_{\min}(\mathbf{A})$ and $\sigma_{\max}(\mathbf{A})$ denote its smallest and largest singular values, with $\kappa(\mathbf{A}) = \sigma_{\max}(\mathbf{A})/\sigma_{\min}(\mathbf{A})$. The notation $\|\cdot\|_2$ denotes the Euclidean norm, $\|\cdot\|_F$ the Frobenius norm, and $\lceil x \rceil$ the smallest integer greater than or equal to x.

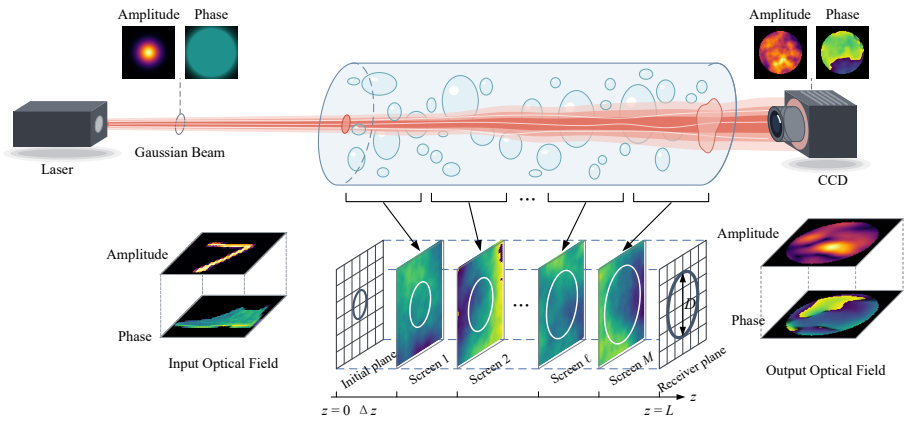


Figure 1: Distributed atmospheric-turbulence propagation with finite-aperture reception. The complex field propagates through M split-step phase screens over a distance L, and the receiver aperture of diameter D truncates the field at the observation plane.

3.1 Atmospheric Turbulence Propagation

Atmospheric turbulence arises from random refractive-index fluctuations that introduce spatially varying phase perturbations into the propagating optical field. In the following, we model propagation under the scalar paraxial approximation and characterize the turbulence strength by the refractive-index structure parameter C_n^2. For a homogeneous path of length L, the plane-wave Fried coherence diameter is

$$r_0 = \left(0.423\, k_0^2 C_n^2 L\right)^{-3/5}, \qquad k_0 = \frac{2\pi}{\lambda}, \tag{1}$$

where k_0 is the optical wavenumber and λ is the wavelength. The ratio D/r_0, with D the receiver aperture diameter, provides a convenient measure of the phase distortion: a larger D/r_0 corresponds to stronger spatial phase variation (Andrews & Phillips, 2005; Schmidt, 2010).

Figure 1 illustrates the distributed-turbulence propagation model considered in this work. The complex field evolves through alternating free-space diffraction and random phase screens over the path, while the receiver aperture truncates the field at the observation plane. To capture distributed turbulence along the propagation path, we use the standard split-step phase-screen model, in which free-space diffraction and random phase modulation alternate sequentially. Let the two-dimensional optical field be vectorized over N spatial samples, so that $\mathbf{U}_0, \mathbf{U}_L \in \mathbb{C}^N$ denote the transmitted field and the field after a distance L, respectively. Then

$$\mathbf{U}_L = \mathbf{H}\mathbf{U}_0, \qquad \mathbf{H} = \mathbf{\Phi}_M \mathbf{T}_{\Delta z} \cdots \mathbf{\Phi}_1 \mathbf{T}_{\Delta z} \in \mathbb{C}^{N\times N}, \tag{2}$$

where M is the number of phase screens, $\Delta z = L/M$ is the screen spacing, $\mathbf{T}_{\Delta z}$ denotes free-space diffraction over Δz, and $\mathbf{\Phi}_\ell = \mathrm{diag}(e^{\mathrm{i}\boldsymbol{\phi}_\ell})$, $\ell = 1, \ldots, M$, is the ℓ-th random phase screen. Here, $\boldsymbol{\phi}_\ell \in \mathbb{R}^N$ denotes the sampled turbulent phase on the computational grid, the exponential is applied entrywise, and $\mathrm{diag}(\cdot)$ forms a diagonal matrix from a vector. For a homogeneous path, all screens share the same C_n^2 and Δz, corresponding to the per-screen coherence diameter $r_{0,\ell} = (0.423\, k_0^2 C_n^2 \Delta z)^{-3/5}$, while their realizations are sampled independently. Free-space diffraction is implemented in the Fourier domain using the paraxial angular-spectrum transfer function, which is given by

$$\mathbf{T}_{\Delta z} = \mathcal{F}^{-1}\mathcal{H}_{\Delta z}\mathcal{F}, \qquad \mathcal{H}_{\Delta z}(f_x, f_y) \simeq e^{\mathrm{i}k_0\Delta z} e^{-\mathrm{i}\pi\lambda\Delta z(f_x^2+f_y^2)}, \tag{3}$$

where $\mathcal{F}$ denotes the unitary two-dimensional discrete Fourier transform, (f_x, f_y) are the spatial frequencies, and the paraxial transfer function $\mathcal{H}_{\Delta z}$ acts as a pointwise multiplication in the frequency domain. Since $|\mathcal{H}_{\Delta z}(f_x, f_y)| = 1$ and $|e^{\mathrm{i}\phi_{\ell,n}}| = 1$ for every spatial sample n, in the absence of absorption and boundary losses, we have

$$\mathbf{T}_{\Delta z}^H \mathbf{T}_{\Delta z} = \mathbf{I}_N, \qquad \mathbf{\Phi}_\ell^H \mathbf{\Phi}_\ell = \mathbf{I}_N, \qquad \mathbf{H}^H\mathbf{H} = \mathbf{I}_N. \tag{4}$$

Thus, although repeated diffraction between random phase screens produces strong phase–intensity coupling, the propagation operator remains unitary in the ideal lossless model. This property belongs

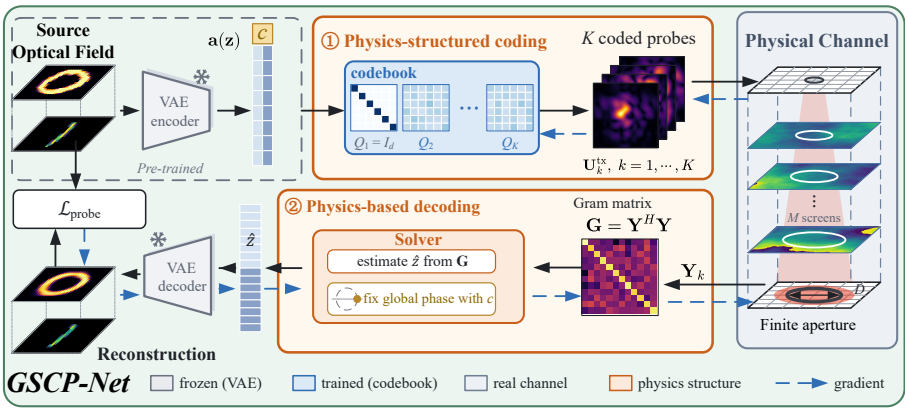


Figure 2: Overview of GSCP-Net. A frozen VAE provides the low-dimensional source representation, a learned unitary codebook generates K coded probes, and a differentiable Gram-equation solver recovers the source from the received-field Gram matrix using the condition-dependent mean effective metric.

to the propagation process and is independent of the subsequent receiver aperture. In the recovery setting considered here, the propagation model and turbulence statistics are specified, whereas the realization-specific phase screens $\mathbf{\Phi}_{1:M} = \{\mathbf{\Phi}_1, \ldots, \mathbf{\Phi}_M\}$ are not provided to the receiver. Rather than estimating these screens or explicitly inverting $\mathbf{H}$, the coded-probe formulation below exploits Eq. 4 as a realization-independent geometric invariant.

To model temporal evolution, we adopt the Taylor frozen-flow hypothesis (Tatarskii, 2016), under which each phase screen translates with a constant transverse velocity $\mathbf{v}_\ell$,

$$\phi_\ell(\mathbf{r}, t) = \phi_\ell(\mathbf{r} - \mathbf{v}_\ell t, 0), \qquad \ell = 1, \ldots, M, \tag{5}$$

and the translated screens are substituted into the split-step model of Eq. 2. Following Bachmann et al. (2023), the transverse velocity is modeled as

$$\mathbf{v}_\ell = \big(\langle V \rangle + \Delta V n_{\ell,1}, \Delta V n_{\ell,2}\big), \qquad n_{\ell,1}, n_{\ell,2} \sim \mathcal{N}(0, 1),$$

where $\langle V \rangle$ is the mean transverse wind speed and ΔV is the standard deviation of the wind components. We characterize the temporal evolution of the turbulence by the atmospheric coherence time

$$t_c = 6.88^{-3/5} \frac{r_0}{\langle V \rangle}, \tag{6}$$

The initial phase-screen realizations are continuously translated rather than independently regenerated at different time instants. The coded-probe acquisition time is assumed to be much shorter than t_c. Hence, the turbulence channel is quasi-static within each coded-probe burst, so that all probes experience the same instantaneous turbulence realization.

3.2 Coded-Probe Forward Model and Gram Geometry

We next construct coded measurements that expose the geometric structure preserved by the propagation model.

Figure 2 summarizes the complete GSCP-Net pipeline.

Let $\mathbf{X}_0 \in \mathbb{C}^N$ denote the complex source field to be recovered, and let $\mathbf{z} = [\mathbf{z}_{\mathrm{re}}^T, \mathbf{z}_{\mathrm{im}}^T]^T \in \mathbb{R}^p$ denote its low-dimensional real representation produced by the source VAE. Specifically, we write $\mathbf{z} = \mathcal{E}(\mathbf{X}_0)$, where $\mathcal{E}$ denotes the deterministic inference map derived from the VAE posterior mean. Here, d denotes the dimension of the transmit subspace, $\mathbf{z}_{\mathrm{re}}, \mathbf{z}_{\mathrm{im}} \in \mathbb{R}^{d-1}$, and $p \triangleq \dim(\mathbf{z}) = 2(d-1)$. We associate this latent representation with the complex coefficient vector

$$\mathbf{a}(\mathbf{z}) = \begin{bmatrix} c \\ \mathbf{z}_{\mathrm{re}} + \mathrm{i}\mathbf{z}_{\mathrm{im}} \end{bmatrix} \in \mathbb{C}^d, \qquad c > 0, \tag{7}$$

where the fixed positive coefficient c serves as a pilot (phase reference), fixing the global-phase ambiguity while controlling the power balance between the reference and source-dependent latent components. The coefficient vector is then embedded into a physical optical field through a semi-unitary transmit basis constructed from training-side propagation simulations and held fixed thereafter. Specifically, $\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_d] \in \mathbb{C}^{N\times d}$, whose columns form an orthonormal basis for the d-dimensional transmit subspace,

$$\mathbf{U}_0 = \mathbf{B}\mathbf{a}(\mathbf{z}) \in \mathbb{C}^N, \qquad \mathbf{B}^H\mathbf{B} = \mathbf{I}_d. \tag{8}$$

Here, $\mathbf{U}_0$ is the latent-encoded optical field transmitted through the turbulence channel, and the source estimate is obtained as $\widehat{\mathbf{X}}_0 = \mathcal{D}(\widehat{\mathbf{z}})$. To generate multiple observations with controlled relative geometry, we apply unitary transformations $\mathbf{Q}_k \in \mathbb{C}^{d\times d}$ to the coefficient vector,

$$\mathbf{U}_k^{\mathrm{tx}} = \mathbf{B}\mathbf{Q}_k\mathbf{a}(\mathbf{z}), \qquad \mathbf{Q}_k^H\mathbf{Q}_k = \mathbf{I}_d, \quad k = 1, \ldots, K, \tag{9}$$

where K is the number of coded probes and $\mathbf{Q}_1 = \mathbf{I}_d$ serves as the reference probe.

After propagation through a common turbulence realization, described by the propagation operator $\mathbf{H}$ in Eq. 2, and reception through a circular aperture of diameter D, the k-th received field is

$$\mathbf{Y}_k = \mathbf{P}_A(D)\mathbf{H}\mathbf{B}\mathbf{Q}_k\mathbf{a}(\mathbf{z}) \in \mathbb{C}^N, \tag{10}$$

where $\mathbf{P}_A(D) \in \mathbb{R}^{N\times N}$ is the diagonal binary mask associated with the receiving aperture, with unit entries inside the aperture and zero entries outside. All K probes within one measurement group are assumed to experience the same instantaneous turbulence realization, as required by the common-channel Gram relation. Following Bachmann et al. (2023), the aperture operator is an orthogonal projector satisfying

$$\mathbf{P}_A^H = \mathbf{P}_A, \qquad \mathbf{P}_A^2 = \mathbf{P}_A, \tag{11}$$

where the dependence on D is omitted for notational simplicity. The limiting case $\mathbf{P}_A = \mathbf{I}_N$ corresponds to ideal lossless FF reception. In this numerical study, we assume coherent access to the complex received fields over the aperture, so that their pairwise inner products can be evaluated. Rather than treating the received probes independently, we characterize their relative geometry through the Gram matrix

$$\mathbf{G} = \mathbf{Y}^H\mathbf{Y} \in \mathbb{C}^{K\times K}, \qquad \mathbf{Y} = [\mathbf{Y}_1, \ldots, \mathbf{Y}_K] \in \mathbb{C}^{N\times K}, \tag{12}$$

with entries $G_{ij} = \mathbf{Y}_i^H\mathbf{Y}_j$. Using $\mathbf{P}_A^H\mathbf{P}_A = \mathbf{P}_A$, the general Gram entry is

$$G_{ij} = \mathbf{a}(\mathbf{z})^H\mathbf{Q}_i^H\mathbf{B}^H\mathbf{H}^H\mathbf{P}_A\mathbf{H}\mathbf{B}\mathbf{Q}_j\mathbf{a}(\mathbf{z}). \tag{13}$$

Under FF reception, $\mathbf{P}_A = \mathbf{I}_N$. Using $\mathbf{H}^H\mathbf{H} = \mathbf{I}_N$ and $\mathbf{B}^H\mathbf{B} = \mathbf{I}_d$, this reduces to

$$G_{ij} = \mathbf{a}(\mathbf{z})^H\mathbf{Q}_i^H\mathbf{Q}_j\mathbf{a}(\mathbf{z}). \tag{14}$$

Therefore, although each received probe depends on the particular turbulence realization, their pairwise Gram geometry is invariant to that realization. Moreover, the Gram matrix is invariant under the common phase rotation $\mathbf{a}(\mathbf{z}) \mapsto e^{\mathrm{i}\theta}\mathbf{a}(\mathbf{z})$ for any $\theta \in \mathbb{R}$, and therefore determines the coefficient vector only up to a global phase. The fixed first entry $c > 0$ of $\mathbf{a}(\mathbf{z})$ removes this ambiguity. By contrast, finite-aperture reception changes this ideal Gram relation. Since $\mathbf{P}_A^H\mathbf{P}_A = \mathbf{P}_A$, the combined propagation–aperture effect within the transmit subspace can be summarized by the effective metric

$$\mathbf{R} = \mathbf{B}^H\mathbf{H}^H\mathbf{P}_A\mathbf{H}\mathbf{B} \in \mathbb{C}^{d\times d}. \tag{15}$$

The corresponding finite-aperture Gram relation becomes

$$G_{ij} = \mathbf{a}(\mathbf{z})^H\mathbf{Q}_i^H\mathbf{R}\mathbf{Q}_j\mathbf{a}(\mathbf{z}). \tag{16}$$

It is also useful to note that, because $\mathbf{P}_A$ is an orthogonal projector, $\mathbf{H}$ is unitary in the ideal lossless model, and $\mathbf{B}^H\mathbf{B} = \mathbf{I}_d$,

$$\mathbf{0} \preceq \mathbf{R} \preceq \mathbf{I}_d, \tag{17}$$

where $\preceq$ denotes the Loewner partial order. Hence, $\mathbf{R}$ is a Hermitian positive-semidefinite contraction with eigenvalues in $[0, 1]$ (Bhatia, 1997). The ideal FF geometry is recovered when $\mathbf{R} = \mathbf{I}_d$. Under finite-aperture truncation, the common propagation operator no longer cancels completely;

instead, its effect on the transmit subspace reduces to the $d \times d$ metric $\mathbf{R}$, rather than the full $N \times N$ operator $\mathbf{H}$.

Finally, denoting the probe codebook by $\mathcal{Q} = \{\mathbf{Q}_k\}_{k=1}^K$, define

$$\mathbf{C}(\mathbf{z}, \mathcal{Q}) = [\mathbf{Q}_1\mathbf{a}(\mathbf{z}), \ldots, \mathbf{Q}_K\mathbf{a}(\mathbf{z})] \in \mathbb{C}^{d\times K}. \tag{18}$$

The complete coded-probe observation model can then be written compactly as

$$\mathbf{G} = \mathbf{C}(\mathbf{z}, \mathcal{Q})^H \mathbf{R} \mathbf{C}(\mathbf{z}, \mathcal{Q}). \tag{19}$$

To distinguish the instantaneous channel from the quantity available to the receiver, let $\mathbf{u} = (C_n^2, D)$ denote the known large-scale condition and let ω index a turbulence realization. We write the corresponding metric as

$$\mathbf{R}(\omega; \mathbf{u}) = \mathbf{B}^H \mathbf{H}(\omega; C_n^2)^H \mathbf{P}_A(D) \mathbf{H}(\omega; C_n^2) \mathbf{B}. \tag{20}$$

The realization-specific metric in Eq. 20 is not supplied to the receiver. Instead, using only training realizations $\{\omega_s\}_{s=1}^{S_{\mathrm{tr}}}$, we form the empirical condition-dependent mean

$$\overline{\mathbf{R}}(\mathbf{u}) = \frac{1}{S_{\mathrm{tr}}} \sum_{s=1}^{S_{\mathrm{tr}}} \mathbf{R}(\omega_s; \mathbf{u}). \tag{21}$$

Because Eq. 21 is a convex average of positive-semidefinite contractions, it also satisfies $\mathbf{0} \preceq \overline{\mathbf{R}}(\mathbf{u}) \preceq \mathbf{I}_d$. For a new realization, define $\Delta\mathbf{R}(\omega; \mathbf{u}) = \mathbf{R}(\omega; \mathbf{u}) - \overline{\mathbf{R}}(\mathbf{u})$. By construction, this fluctuation has zero empirical mean over the training subset at each condition. The measured Gram matrix therefore contains the realization-dependent perturbation

$$\mathbf{G}_{\mathrm{obs}} = \mathbf{C}^H \overline{\mathbf{R}}(\mathbf{u}) \mathbf{C} + \mathbf{C}^H \Delta\mathbf{R}(\omega; \mathbf{u}) \mathbf{C}, \tag{22}$$

whereas the receiver uses only $\overline{\mathbf{R}}(\mathbf{u})$. Probe optimization therefore accounts for realization-dependent metric variations without supplying the instantaneous metric to the receiver. The experiments below evaluate the resulting robustness on independently generated turbulence realizations.

3.3 Identifiability and Conditioning Analysis

To characterize the observability of the latent source under the same equations used by the physical receiver, we retain the complex off-diagonal Gram entries together with the reference-probe power G_{11}. For a fixed effective metric $\mathbf{R}$, define

$$\mathbf{g}_{\mathbf{R}}(\mathbf{z}) = [G_{11}, \{\Re G_{ij}, \Im G_{ij}\}_{i<j}]^T \in \mathbb{R}^m, \qquad m = 1 + K(K-1). \tag{23}$$

Under FF reception, $\mathbf{R} = \mathbf{I}_d$ and Eq. 23 contains all independent real Gram coordinates because all diagonal entries are equal. Under finite-aperture reception, the remaining diagonal powers are generally distinct, but they are not used by the receiver considered here. The corresponding Jacobian is

$$\mathbf{J}_{\mathbf{R}} = \frac{\partial \mathbf{g}_{\mathbf{R}}(\mathbf{z})}{\partial \mathbf{z}^T} \in \mathbb{R}^{m\times p}. \tag{24}$$

For FF reception we write $\mathbf{J}_{\mathrm{FF}} = \mathbf{J}_{\mathbf{I}_d}$. For brevity, we write $\mathbf{J} = \mathbf{J}_{\mathbf{R}}$ below.

In the local linearized model, all p latent directions are observable if and only if

$$\mathrm{rank}(\mathbf{J}) = p. \tag{25}$$

This rank condition also yields a simple counting bound on the number of coded probes. Since Eq. 23 contains $m = 1 + K(K-1)$ real equations, a necessary condition for $\mathbf{J}_{\mathbf{R}}$ to have full column rank is

$$1 + K(K-1) \geq p, \tag{26}$$

or equivalently,

$$K \geq \left\lceil \frac{1 + \sqrt{4p-3}}{2} \right\rceil. \tag{27}$$

Full rank alone does not guarantee stable inversion. We therefore further characterize the local recovery geometry using the minimum singular value $\sigma_{\min}(\mathbf{J})$ and the condition number $\kappa(\mathbf{J})$. Larger $\sigma_{\min}$ and smaller κ indicate better conditioning. Detailed derivations for both FF and finite-aperture reception are provided in Appendix A.

3.4 PHYSICS-BASED RECEIVER AND TRAINING PROCEDURE

Given $\mathbf{G}_{\text{obs}}$, $\mathcal{Q}$, and $\mathbf{u} = (C_n^2, D)$, the receiver estimates the latent source using $\overline{\mathbf{R}}(\mathbf{u})$ through the physics-based inverse problem defined by the Gram model. The receiver itself has no trainable parameters. The probes are parameterized as

$$\mathbf{Q}_1 = \mathbf{I}_d, \qquad \mathbf{Q}_k = \mathbf{Q}_k^0 \exp(\mathrm{i}\mathbf{A}_k), \qquad \mathbf{A}_k = \mathbf{A}_k^H, \quad k = 2, \ldots, K, \tag{28}$$

where $\mathbf{Q}_k^0$ is the initial unitary probe, $\mathbf{A}_k$ is a trainable Hermitian matrix, and $\exp(\cdot)$ denotes the matrix exponential. This parameterization preserves exact unitarity. During probe optimization, observations are generated with sampled $\mathbf{R}(\omega; \mathbf{u})$, whereas the receiver uses only $\overline{\mathbf{R}}(\mathbf{u})$. Back-propagation through the differentiable Gram-equation solver updates only $\{\mathbf{A}_k\}_{k=2}^K$. The complete construction and training procedure is summarized in Algorithms 1 and 2 in Appendix B.

4 EXPERIMENTS

Source fields are constructed from independently sampled MNIST amplitude and phase maps A and P, normalized to $[0, 1]$. We set $A_{\text{phys}} = 0.15 + 0.85A$ and $\phi_{\text{phys}} = \frac{\pi}{2}(2P - 1)$, followed by a fixed super-Gaussian envelope and unit-energy normalization. A pretrained source VAE maps each two-channel 28×28 field to $p = 64$ real latent variables, corresponding to $d = 33$, with $c \approx \sqrt{\mathbb{E}\|\mathbf{z}\|_2^2}$. The VAE is frozen before probe optimization.

The transmit basis $\mathbf{B}$ is constructed offline from a 64-dimensional beam-localized candidate basis $\mathbf{B}_{\text{cand}} \in \mathbb{C}^{N \times 64}$. At $C_n^2 = 10^{-13}\,\mathrm{m}^{-2/3}$ and $D = 220$ mm, the mean candidate-space metric is estimated from 32 training-side realizations; its leading $d = 33$ modes are mapped through $\mathbf{B}_{\text{cand}}$ and QR-orthonormalized. The resulting basis is fixed thereafter.

Propagation is simulated over 2 km at $\lambda = 1550$ nm on a 512×512 grid with 2.5-mm sampling and 100 von Kármán phase screens ($L_0 = 50$ m, $l_0 = 20$ μm, three subharmonic levels), without realization-wise RMS rescaling. The beam has a 160-mm $1/e^2$ intensity diameter.

We evaluate six turbulence strengths, $C_n^2 \in \{10^{-15}, 3 \times 10^{-15}, 10^{-14}, 3 \times 10^{-14}, 10^{-13}, 3 \times 10^{-13}\}\,\mathrm{m}^{-2/3}$, and seven receiver apertures, $D \in \{140, 180, 220, 260, 320, 400, 480\}$ mm, yielding 42 turbulence–aperture conditions. Mean-metric estimation and probe optimization use training-side turbulence realizations, while evaluation is performed on an independent test set with the same realizations reused across methods for paired comparison. The corresponding physical conditions are summarized in Table 11.

Unless otherwise stated, $K = 12$ coded probes are used. Complex-field recovery is evaluated using the normalized squared field overlap

$$F\left(\widehat{\mathbf{X}}_0, \mathbf{X}_0\right) = \frac{\left|\widehat{\mathbf{X}}_0^H \mathbf{X}_0\right|^2}{\|\widehat{\mathbf{X}}_0\|_2^2 \, \|\mathbf{X}_0\|_2^2}, \tag{29}$$

following Kupianskyi et al. (2023).

4.1 COMPLEX-FIELD RECOVERY

Table 1: Complex-field recovery fidelity under representative turbulence–aperture conditions. Source fields are constructed from independently sampled MNIST amplitude and phase maps

C_n^2 ($\mathrm{m}^{-2/3}$)	1×10^{-14}			1×10^{-13}			3×10^{-13}		
D (mm)	140	220	480	140	220	480	140	220	480
U-Net (Tan et al., 2022)	0.312	0.384	0.415	0.105	0.097	0.096	0.117	0.099	0.088
NAFNet (Chen et al., 2022)	0.305	0.376	0.397	0.084	0.080	0.078	0.085	0.078	0.073
GSCP-Net ($K = 12$)	0.843	**0.860**	**0.861**	0.669	0.823	**0.861**	0.542	0.720	0.853
GSCP-Net ($K = 15$)	**0.850**	**0.860**	**0.861**	**0.742**	**0.839**	**0.861**	**0.660**	**0.778**	**0.857**
Source VAE reconstruction	0.861	0.861	0.861	0.861	0.861	0.861	0.861	0.861	0.861

We compare GSCP-Net with representative end-to-end baselines on the same paired test sources and turbulence realizations. The frozen VAE alone achieves a mean reconstruction fidelity of 0.861 between $\mathcal{D}(\mathcal{E}(\mathbf{X}_0))$ and $\mathbf{X}_0$. Figure 3 complements the aggregate results in Table 1 with representative amplitude–phase reconstructions. The end-to-end baselines exhibit increasing distortion as turbulence strengthens and the aperture becomes more restrictive, whereas GSCP-Net preserves the source structure more consistently across the illustrated conditions.

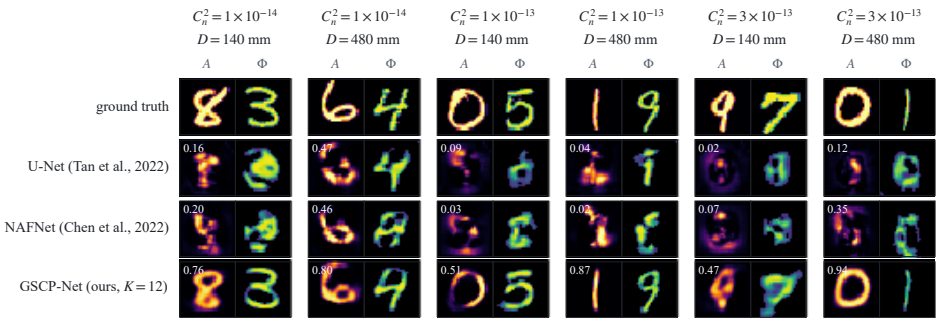


Figure 3: Representative amplitude–phase reconstructions under different turbulence strengths and receiver apertures. Rows compare the ground truth, U-Net (Tan et al., 2022), NAFNet (Chen et al., 2022), and GSCP-Net ($K = 12$); displayed values are per-example field-overlap fidelities.

Table 1 shows that GSCP-Net consistently outperforms both baselines, with the largest margins under stronger turbulence and smaller apertures. Over all 42 conditions, the mean fidelities are 0.830 and 0.843 for $K = 12$ and $K = 15$, with worst-case values of 0.542 and 0.660. At larger apertures, fidelity approaches the 0.861 VAE reconstruction level even under strong turbulence. Complete 42-condition and PSNR/SSIM results are reported in Appendix C; additional qualitative sweeps are shown in Figs. 4–6. GSCP-Net also learns only about 1.2×10^4 channel-dependent parameters, compared with 7.8×10^6 for U-Net and 4.0×10^6 for NAFNet, and trains in about 40 min instead of 2.9–4.0 h on the same GPU (Table 12).

4.2 Probe Design and Number of Probes

We isolate the contribution of probe design and codebook size. For $p = 64$, Sec. 3.3 gives the necessary bound $K \geq 9$; we therefore compare learned $K = 9, 12, 15$ codebooks and the fixed $K = 12$ initialization $\{\mathbf{Q}_k^0\}_{k=1}^K$.

Table 2: Complex-field recovery fidelity for probe-design and probe-number ablations.

C_n^2 ($\mathrm{m}^{-2/3}$)	10^{-13}			3×10^{-13}			all 42 conditions	
D (mm)	140	220	480	140	220	480	mean	worst
Learned probes $\mathbf{Q}_k$ ($K = 12$)	0.669	0.823	0.861	0.542	0.720	0.853	0.830	0.542
Fixed initial probes $\mathbf{Q}_k^0$ ($K = 12$)	0.427	0.748	0.861	0.302	0.511	0.837	0.786	0.302
Learned probes, $K = 9$	0.289	0.650	0.861	0.189	0.346	0.802	0.750	0.189
Learned probes, $K = 15$	0.742	0.839	0.861	0.660	0.778	0.857	0.843	0.660

Table 2 shows that probe optimization is most beneficial at small apertures. At $C_n^2 = 3 \times 10^{-13}\,\mathrm{m}^{-2/3}$ and $D = 140$ mm, the $K = 12$ fidelity rises from 0.302 to 0.542, with smaller gains as aperture truncation weakens. Although $K = 9$ satisfies the counting bound, it degrades under strong finite-aperture truncation; increasing to $K = 15$ raises the worst-case fidelity from 0.542 to 0.660. The conditioning analysis shows the same trend: at $K = 9$, probe learning increases the median $\sigma_{\min}(\mathbf{J})$ from 0.58 to 0.93 and reduces the median $\kappa(\mathbf{J})$ from 22.8 to 14.0, whereas at $K = 15$ both fixed and learned codebooks have median $\kappa = 3.3$. Thus, additional probes improve conditioning while reducing the marginal benefit of probe optimization; see Table 7 and Fig. 7.

4.3 Robustness to Dynamic Turbulence

We finally test dynamic turbulence using the Taylor frozen-flow model in Eq. 5, with $\langle V \rangle = 3$ m/s and $\Delta V = 1$ m/s. Each condition uses 12 turbulence realizations and 64 test sources. The learned $K = 12$ codebook and receiver remain fixed, using the same training-derived $\overline{\mathbf{R}}(C_n^2, D)$ at all times; the instantaneous $\mathbf{R}(t)$ is reported only for diagnosis.

Table 3: Complex-field recovery fidelity under dynamically evolving turbulence, together with the relative change in the effective metric at $20t_c$.

C_n^2 ($\mathrm{m}^{-2/3}$)	1×10^{-13}			3×10^{-13}		
D (mm)	140	220	480	140	220	480
$t = 0$ (static)	0.682	0.819	0.858	0.555	0.745	0.855
$t = 5\,t_c$	0.689	0.822	0.858	0.561	0.748	0.855
$t = 10\,t_c$	0.698	0.825	0.858	0.556	0.751	0.855
$t = 20\,t_c$	0.682	0.824	0.858	0.546	0.747	0.853
$\|\mathbf{R}(20t_c) - \mathbf{R}(0)\|_F / \|\mathbf{R}(0)\|_F$	0.83	0.42	0.02	0.91	0.60	0.10

Table 3 shows nearly unchanged fidelity over $20\,t_c$. At $C_n^2 = 3 \times 10^{-13}\,\mathrm{m}^{-2/3}$ and $D = 140$ mm, for example, $\|\mathbf{R}(20t_c) - \mathbf{R}(0)\|_F / \|\mathbf{R}(0)\|_F = 0.91$ while fidelity changes only from 0.555 to 0.546. Further dynamic results are shown in Table 8 and Fig. 8.

Additional off-grid and cross-dataset results are given in Tables 9 and 10, with qualitative examples in Fig. 9; implementation details are in Table 12.

5 Discussion

Under FF reception, the coded-probe Gram relation is exactly invariant to the turbulence realization in the ideal lossless model; finite apertures instead introduce the low-dimensional metric $\mathbf{R}(\omega; \mathbf{u})$. The receiver uses only the training-derived mean $\overline{\mathbf{R}}(\mathbf{u})$, while probe optimization accounts for realization-specific variations. The ablations show that learned and additional probes are most beneficial under strong aperture truncation. Static and dynamic tests use realizations independent of those used to construct $\overline{\mathbf{R}}$ and optimize the probes. Supplementary off-grid and cross-dataset results further support robustness beyond the main benchmark. Unmodeled amplitude or absorption effects, additional propagation mismatch, and experimental hardware are left for future work.

6 Conclusion

We introduced GSCP-Net for complex-field recovery through atmospheric turbulence under finite-aperture reception, combining a low-dimensional source representation, exactly unitary coded probes, Gram-structured physics, and physics-based inversion. Under finite apertures, the ideal FF Gram invariance breaks, and the joint effect of propagation and truncation reduces to the low-dimensional effective metric $\mathbf{R}$. Across 42 turbulence–aperture conditions over a 2-km distributed-turbulence channel, GSCP-Net substantially outperformed representative end-to-end baselines, with the largest gains under strong turbulence and small apertures, while learning orders of magnitude fewer channel-dependent parameters and requiring only a fraction of their training time. Probe optimization and additional probes improved worst-case recovery, while the same learned probes and fixed receiver maintained nearly unchanged fidelity under dynamically evolving turbulence without instantaneous channel tracking. These results establish Gram-structured coded probing as a lightweight, physics-guided approach to robust complex-field recovery across a broad range of turbulence strengths and receiver apertures.

AI Use Statement

Generative AI tools were used to assist with language editing, presentation, and limited drafting and organizational support. All scientific content, mathematical derivations, numerical experiments,

reported results, and conclusions were reviewed and verified by the authors, who take full responsibility for the final manuscript.

Reproducibility Statement

The propagation model, source representation, transmit-basis construction, training/test split, probe parameterization, physics-based receiver, optimization procedure, and evaluation protocol are specified in the main text and Appendix B. All reported test turbulence realizations are independent of those used to construct the mean metrics and optimize the probes.

References

Larry C. Andrews and Ronald L. Phillips. *Laser Beam Propagation through Random Media*. SPIE Press, Bellingham, WA, 2nd edition, 2005.

Jorge Bacca, Tatiana Gelvez, and Henry Arguello. Deep coded aperture design: An end-to-end approach for computational imaging tasks. *IEEE Transactions on Computational Imaging*, 7: 1148–1160, 2021. doi: 10.1109/TCI.2021.3122285.

David Bachmann, Mathieu Isoard, Vyacheslav Shatokhin, Giacomo Sorelli, Nicolas Treps, and Andreas Buchleitner. Highly transmitting modes of light in dynamic atmospheric turbulence. *Physical review letters*, 130(7):073801, 2023.

Rajendra Bhatia. *Matrix Analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer, New York, 1997.

Julie Chang and Gordon Wetzstein. Deep optics for monocular depth estimation and 3D object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10193–10202, 2019.

Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *European conference on computer vision*, pp. 17–33. Springer, 2022.

Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pp. 2921–2926. IEEE, 2017.

Bertrand M Hochwald and Thomas L Marzetta. Unitary space-time modulation for multiple-antenna communications in rayleigh flat fading. *IEEE Transactions on Information Theory*, 46(2):543–564, 2000.

Qi Jia, Rui Feng, Bojian Shi, Fangkui Sun, Yanxia Zhang, Hang Li, Xiaoxin Li, Yongyin Cao, Jian Wang, and Weiqiang Ding. Compensating the distorted oam beams with near zero time delay. *Applied Physics Letters*, 121(1), 2022.

Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.

Hlib Kupianskyi, Simon AR Horsley, and David B Phillips. High-dimensional spatial mode sorting and optical circuit design using multi-plane light conversion. *Apl Photonics*, 8(2), 2023.

Junmin Liu, Peipei Wang, Xiaoke Zhang, Yanliang He, Xinxing Zhou, Huapeng Ye, Ying Li, Shixiang Xu, Shuqing Chen, and Dianyuan Fan. Deep learning based atmospheric turbulence compensation for orbital angular momentum beam distortion and communication. *Optics express*, 27 (12):16671–16688, 2019.

Casey J Pellizzari, Mark F Spencer, and Charles A Bouman. Coherent plug-and-play: digital holographic imaging through atmospheric turbulence using model-based iterative reconstruction and convolutional neural networks. *IEEE Transactions on Computational Imaging*, 6:1607–1621, 2020.

Cameron J Radosevich, Casey J Pellizzari, Samuel Horst, and Mark F Spencer. Imaging through deep turbulence using single-shot digital holography data. *Optics Express*, 28(13):19390–19401, 2020.

Jason D. Schmidt. *Numerical Simulation of Optical Wave Propagation with Examples in MATLAB*. SPIE Press, Bellingham, WA, 2010.

Yehong Tan, Xuanyu Hu, and Jian Wang. Complex amplitude field reconstruction in atmospheric turbulence based on deep learning. *Optics Express*, 30(8):13070–13078, 2022.

V. I. Tatarskii. *Wave Propagation in a Turbulent Medium*. Dover Publications, Mineola, NY, 2016.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Wenjie Xiong, Peipei Wang, Menglong Cheng, Junmin Liu, Yanliang He, Xinxing Zhou, Jiangnan Xiao, Ying Li, Shuqing Chen, and Dianyuan Fan. Convolutional neural network based atmospheric turbulence compensation for optical orbital angular momentum multiplexing. *Journal of Lightwave Technology*, 38(7):1712–1721, 2020.

A Detailed Identifiability and Conditioning Analysis

For completeness, this appendix provides the detailed derivation of the identifiability and conditioning analysis summarized in Sec. 3.3. We first consider FF reception and then extend the analysis to finite-aperture reception.

Under FF reception, $\mathbf{R} = \mathbf{I}_d$, and Eq. 19 reduces to

$$\mathbf{G} = \mathbf{C}(\mathbf{z}, \mathcal{Q})^H \mathbf{C}(\mathbf{z}, \mathcal{Q}). \tag{30}$$

Since every $\mathbf{Q}_k$ is unitary,

$$G_{kk} = \mathbf{a}(\mathbf{z})^H \mathbf{a}(\mathbf{z}), \qquad k = 1, \dots, K, \tag{31}$$

so all diagonal entries of $\mathbf{G}$ are identical. Because $\mathbf{G}$ is Hermitian, the FF Gram feature vector contains at most

$$m_{\mathrm{FF}} = 1 + K(K-1), \tag{32}$$

independent real coordinates. With the feature ordering of Eq. 23, we write $\mathbf{g}_{\mathrm{FF}} = \mathbf{g}_{\mathbf{I}_d}$.

The corresponding FF Jacobian is

$$\mathbf{J}_{\mathrm{FF}} = \frac{\partial \mathbf{g}_{\mathrm{FF}}(\mathbf{z})}{\partial \mathbf{z}^T} \in \mathbb{R}^{m_{\mathrm{FF}} \times p}. \tag{33}$$

For finite-aperture reception, conditional on a fixed effective metric $\mathbf{R}$, the Gram matrix is given by Eq. 19, with $\mathbf{0} \preceq \mathbf{R} \preceq \mathbf{I}_d$. The receiver uses the same $m = 1 + K(K-1)$ real equations as in Eq. 23: all complex off-diagonal entries together with the reference power G_{11}. Therefore,

$$\mathbf{J}_{\mathbf{R}} = \frac{\partial \mathbf{g}_{\mathbf{R}}(\mathbf{z})}{\partial \mathbf{z}^T} \in \mathbb{R}^{m \times p}. \tag{34}$$

The remaining $K-1$ diagonal powers are physically measurable under finite-aperture reception but are not used by the main receiver.

Both Jacobians follow from the same first-order variation. For

$$\delta\mathbf{z} = [\delta\mathbf{z}_{\mathrm{re}}^T, \delta\mathbf{z}_{\mathrm{im}}^T]^T,$$

Eq. 7 gives

$$\delta\mathbf{a} = \begin{bmatrix} 0 \\ \delta\mathbf{z}_{\mathrm{re}} + \mathrm{i}\,\delta\mathbf{z}_{\mathrm{im}} \end{bmatrix}. \tag{35}$$

Accordingly,

$$\delta\mathbf{C} = [\mathbf{Q}_1\delta\mathbf{a}, \dots, \mathbf{Q}_K\delta\mathbf{a}],$$

and, for fixed $\mathbf{R}$,

$$\delta\mathbf{G} = \delta\mathbf{C}^H\mathbf{R}\mathbf{C}(\mathbf{z}, \mathcal{Q}) + \mathbf{C}(\mathbf{z}, \mathcal{Q})^H\mathbf{R}\delta\mathbf{C}. \tag{36}$$

Hence,

$$\delta\mathbf{g} = \mathbf{J}\delta\mathbf{z}, \tag{37}$$

where $\mathbf{J} = \mathbf{J}_{\mathbf{R}}$ and $\mathbf{J}_{\mathrm{FF}} = \mathbf{J}_{\mathbf{I}_d}$ in the FF case.

Proposition A.1 (First-order identifiability)**.** *All p latent directions are observable in the local linearized model if and only if*

$$\mathrm{rank}(\mathbf{J}) = p. \tag{38}$$

Proof. If $\mathrm{rank}(\mathbf{J}) = p$, then $\ker(\mathbf{J}) = \{\mathbf{0}\}$, so every nonzero $\delta\mathbf{z}$ produces a nonzero first-order change in the Gram features.

If $\mathrm{rank}(\mathbf{J}) < p$, then

$$\dim\ker(\mathbf{J}) = p - \mathrm{rank}(\mathbf{J}) > 0, \tag{39}$$

so there exists a nonzero $\delta\mathbf{z}$ satisfying

$$\mathbf{J}\delta\mathbf{z} = \mathbf{0}.$$

□

For the receiver feature set in Eq. 23, both FF and finite-aperture Jacobians have $m = 1 + K(K-1)$ rows. Hence, a necessary counting condition is

$$1 + K(K-1) \geq p. \tag{40}$$

Equivalently,

$$K \geq \left\lceil \frac{1 + \sqrt{4p-3}}{2} \right\rceil. \tag{41}$$

This counting condition is only necessary. Finite-aperture truncation changes the Jacobian through $\mathbf{R}$ and can strongly degrade its singular-value spectrum even when full column rank is preserved.

Finally,

$$\sigma_{\min}(\mathbf{J})\|\delta\mathbf{z}\|_2 \leq \|\mathbf{J}\delta\mathbf{z}\|_2 \leq \sigma_{\max}(\mathbf{J})\|\delta\mathbf{z}\|_2, \tag{42}$$

where $\sigma_{\min}(\mathbf{J})$ characterizes the weakest observable latent direction. The condition number is

$$\kappa(\mathbf{J}) = \frac{\sigma_{\max}(\mathbf{J})}{\sigma_{\min}(\mathbf{J})}, \tag{43}$$

and it characterizes the anisotropy of the local inverse problem.

B GSCP-NET IMPLEMENTATION AND OPTIMIZATION DETAILS

This appendix summarizes the implementation used to obtain the reported results. The key separation is between (i) the realization-specific metric used to generate each observation and (ii) the training-derived condition-dependent mean metric used by the receiver.

B.1 STAGE A: STATIC CONSTRUCTION AND MEAN METRIC

Algorithm 1 Stage A: Static Construction and Training-Derived Mean Metric

Require: Source training set, candidate transmit basis $\mathbf{B}_{\mathrm{cand}}$, design condition $\mathbf{u}_{\mathrm{des}} = (C_{n,\mathrm{des}}^2, D_{\mathrm{des}})$, and condition grid $\mathcal{S}_{C_n^2} \times \mathcal{S}_D$

Ensure: Frozen source VAE, fixed semi-unitary basis $\mathbf{B}$, training metric library $\mathcal{R}_{\mathrm{tr}}$, training-derived mean metric $\overline{\mathbf{R}}(C_n^2, D)$, and initial unitary codebook $\mathcal{Q}^0$

1: Train the source VAE and freeze its encoder and decoder.

2: At the design condition, compute the candidate-space metrics

$$\mathbf{R}_{\mathrm{cand}}^{(s)} = \mathbf{B}_{\mathrm{cand}}^H \mathbf{H}^{(s)H} \mathbf{P}_A(D_{\mathrm{des}}) \mathbf{H}^{(s)} \mathbf{B}_{\mathrm{cand}},$$

form their empirical mean, retain its leading d eigenvectors, map them through $\mathbf{B}_{\mathrm{cand}}$, and orthonormalize the result to obtain $\mathbf{B}$ with $\mathbf{B}^H\mathbf{B} = \mathbf{I}_d$.

3: **for** each condition-grid condition $\mathbf{u} = (C_n^2, D)$ **do**

4: Generate independent turbulence realizations and compute

$$\mathbf{R}^{(s)}(\mathbf{u}) = \mathbf{B}^H \mathbf{H}^{(s)H} \mathbf{P}_A(D) \mathbf{H}^{(s)} \mathbf{B}.$$

5: Let $\mathcal{I}_{\mathrm{tr}}$ denote the training-realization index set, with $|\mathcal{I}_{\mathrm{tr}}| = S_{\mathrm{tr}}$. Using only these training realizations, form

$$\overline{\mathbf{R}}(\mathbf{u}) = \frac{1}{S_{\mathrm{tr}}} \sum_{s \in \mathcal{I}_{\mathrm{tr}}} \mathbf{R}^{(s)}(\mathbf{u}),$$

and store the corresponding realization-specific metrics in $\mathcal{R}_{\mathrm{tr}}(\mathbf{u})$.

6: **end for**

7: Set $\mathbf{Q}_1^0 = \mathbf{I}_d$ and initialize $\mathbf{Q}_2^0, \ldots, \mathbf{Q}_K^0$ as unitary matrices.

8: **return** Frozen source VAE, $\mathbf{B}$, $\mathcal{R}_{\mathrm{tr}}$, $\overline{\mathbf{R}}$, and $\mathcal{Q}^0$.

For the reported experiments, the candidate basis has dimension 64 and is designed at $C_n^2 = 10^{-13}\,\mathrm{m}^{-2/3}$ and $D = 220$ mm using 32 independent realizations. For each of the 42 turbulence–aperture conditions, 96 realization-specific metrics are generated; 64 are used for mean-metric estimation and probe optimization, while the remaining 32 are held out from optimization. The benchmark results in the main text use a separately generated test set.

B.2 Stage B: Stochastic Probe Optimization

During training, each Gram measurement is generated using a realization-specific metric, while the receiver uses only the corresponding training-derived mean metric.

Algorithm 2 Stage B: Stochastic Probe Optimization

Require: Frozen source VAE and $\mathbf{B}$, training metric library $\mathcal{R}_{\mathrm{tr}}$, training-derived mean metric $\overline{\mathbf{R}}(C_n^2, D)$, and initial unitary codebook $\mathcal{Q}^0$

Ensure: Learned unitary codebook $\mathcal{Q}^\star$

1: Set $\mathbf{Q}_1 = \mathbf{I}_d$ and parameterize $\mathbf{Q}_k = \mathbf{Q}_k^0 \exp(\mathrm{i}\mathbf{A}_k)$, $\mathbf{A}_k = \mathbf{A}_k^H$, for $k = 2, \ldots, K$.

2: **for** each training iteration **do**

3: Sample a source $\mathbf{X}_0$ and a condition-grid condition $\mathbf{u} = (C_n^2, D)$, then compute $\mathbf{z} = \mathcal{E}(\mathbf{X}_0)$.

4: Draw $\mathbf{R}_{\mathrm{true}} \in \mathcal{R}_{\mathrm{tr}}(\mathbf{u})$ and generate

$$\mathbf{G}_{\mathrm{obs}} = \mathbf{C}(\mathbf{z}, \mathcal{Q})^H \mathbf{R}_{\mathrm{true}} \mathbf{C}(\mathbf{z}, \mathcal{Q}).$$

5: Set $\mathbf{R}_{\mathrm{rec}} = \overline{\mathbf{R}}(\mathbf{u})$ and recover

$$\widehat{\mathbf{z}} = \arg\min_{\mathbf{z}} \frac{1}{2} \left\| \mathbf{g}_{\mathbf{R}_{\mathrm{rec}}}(\mathbf{z}) - \mathbf{g}_{\mathrm{obs}} \right\|_2^2 .$$

6: Decode $\widehat{\mathbf{X}}_0 = \mathcal{D}(\widehat{\mathbf{z}})$ and minimize

$$\mathcal{L}_{\mathrm{probe}} = 1 - F(\widehat{\mathbf{X}}_0, \mathbf{X}_0) + 0.10\, \mathcal{L}_{\mathrm{map}}.$$

7: Backpropagate through the differentiable physics-based receiver and update only $\{\mathbf{A}_k\}_{k=2}^K$.

8: **end for**

9: **return** $\mathcal{Q}^\star$.

Here, $\mathbf{g}_{\mathrm{obs}}$ is constructed from $\mathbf{G}_{\mathrm{obs}}$ using the Gram features defined in Eq. 23, and $\mathcal{L}_{\mathrm{map}}$ denotes the mean-squared error between the decoded and target two-channel amplitude–phase maps.

C Additional Experimental Results

C.1 Complete Complex-Field Recovery Results

Table 4: Complex-field recovery fidelity over all 42 turbulence–aperture conditions, with A and P independently sampled from MNIST. Complete results corresponding to Table 1.

C_n^2 ($\mathrm{m}^{-2/3}$) \ D (mm)	140	180	220	260	320	400	480
U-Net (Tan et al., 2022)							
1×10^{-15}	0.514	0.565	0.600	0.624	0.640	0.642	0.640
3×10^{-15}	0.446	0.498	0.533	0.556	0.571	0.573	0.570
1×10^{-14}	0.312	0.355	0.384	0.402	0.415	0.417	0.415
3×10^{-14}	0.162	0.175	0.186	0.192	0.197	0.199	0.199
1×10^{-13}	0.105	0.100	0.097	0.095	0.094	0.095	0.096
3×10^{-13}	0.117	0.103	0.099	0.095	0.090	0.089	0.088
NAFNet (Chen et al., 2022)							
1×10^{-15}	0.529	0.580	0.611	0.631	0.643	0.646	0.645
3×10^{-15}	0.454	0.506	0.539	0.558	0.569	0.571	0.571
1×10^{-14}	0.305	0.352	0.376	0.390	0.396	0.397	0.397
3×10^{-14}	0.138	0.148	0.156	0.160	0.162	0.163	0.164
1×10^{-13}	0.084	0.082	0.080	0.080	0.079	0.078	0.078
3×10^{-13}	0.085	0.080	0.078	0.075	0.074	0.073	0.073
GSCP-Net ($K = 12$)							
1×10^{-15}	0.859	0.861	0.861	0.861	0.861	0.861	0.861
3×10^{-15}	0.855	0.859	0.861	0.861	0.861	0.861	0.861
1×10^{-14}	0.843	0.856	0.860	0.861	0.861	0.861	0.861
3×10^{-14}	0.812	0.845	0.856	0.859	0.861	0.861	0.861
1×10^{-13}	0.669	0.772	0.823	0.846	0.857	0.861	0.861
3×10^{-13}	0.542	0.647	0.720	0.767	0.810	0.840	0.853
GSCP-Net ($K = 15$)							
1×10^{-15}	0.860	0.861	0.861	0.861	0.861	0.861	0.861
3×10^{-15}	0.857	0.860	0.861	0.861	0.861	0.861	0.861
1×10^{-14}	0.850	0.858	0.860	0.861	0.861	0.861	0.861
3×10^{-14}	0.831	0.851	0.858	0.860	0.861	0.861	0.861
1×10^{-13}	0.742	0.809	0.839	0.852	0.859	0.861	0.861
3×10^{-13}	0.660	0.732	0.778	0.807	0.832	0.849	0.857



Figure 4: Expanded qualitative comparison across the six turbulence strengths for (a) $D = 140$ mm and (b) $D = 480$ mm. Amplitude A and phase Φ are shown for the ground truth, U-Net, NAFNet, and GSCP-Net.

Table 5: PSNR and SSIM of the recovered optical fields under representative turbulence–aperture conditions.

C_n^2 ($m^{-2/3}$)	1×10^{-14}			1×10^{-13}			3×10^{-13}			all 42 cells	
D (mm)	140	220	480	140	220	480	140	220	480	mean	worst
PSNR of the amplitude map (dB)											
U-Net (Tan et al., 2022)	11.78	12.45	12.80	10.54	10.53	10.47	10.56	10.48	10.41	12.10	10.41
NAFNet (Chen et al., 2022)	11.60	12.24	12.47	10.25	10.22	10.13	10.22	10.16	10.08	11.89	10.08
GSCP-Net ($K = 12$)	18.16	18.60	**18.65**	15.24	17.74	**18.64**	13.67	15.89	18.42	18.05	13.67
GSCP-Net ($K = 15$)	**18.33**	**18.62**	**18.65**	**16.30**	**18.08**	**18.65**	**15.00**	**16.79**	**18.51**	**18.24**	**15.00**
Source VAE reconstruction	18.65	18.65	18.65	18.65	18.65	18.65	18.65	18.65	18.65	18.65	18.65
PSNR of the phase map (dB)											
U-Net (Tan et al., 2022)	12.40	13.12	13.41	10.83	10.70	10.61	11.06	10.87	10.73	12.70	10.61
NAFNet (Chen et al., 2022)	12.43	13.12	13.32	10.55	10.43	10.38	10.66	10.51	10.37	12.59	10.37
GSCP-Net ($K = 12$)	20.51	21.02	**21.07**	17.20	20.02	**21.07**	15.40	17.94	20.82	20.39	15.40
GSCP-Net ($K = 15$)	**20.70**	**21.04**	**21.07**	**18.34**	**20.42**	**21.07**	**16.98**	**18.98**	**20.92**	**20.61**	**16.98**
Source VAE reconstruction	21.07	21.07	21.07	21.07	21.07	21.07	21.07	21.07	21.07	21.07	21.07
SSIM of the amplitude map											
U-Net (Tan et al., 2022)	0.251	0.327	0.371	0.066	0.077	0.092	0.042	0.044	0.054	0.259	0.042
NAFNet (Chen et al., 2022)	0.248	0.367	0.413	0.067	0.090	0.101	0.046	0.062	0.071	0.294	0.046
GSCP-Net ($K = 12$)	0.773	0.789	**0.790**	0.627	0.755	**0.790**	0.529	0.662	0.781	0.763	0.529
GSCP-Net ($K = 15$)	**0.779**	**0.789**	**0.790**	**0.685**	**0.769**	**0.790**	**0.614**	**0.709**	**0.785**	**0.773**	**0.614**
Source VAE reconstruction	0.790	0.790	0.790	0.790	0.790	0.790	0.790	0.790	0.790	0.790	0.790
SSIM of the phase map											
U-Net (Tan et al., 2022)	0.328	0.401	0.432	0.140	0.146	0.138	0.128	0.130	0.117	0.336	0.117
NAFNet (Chen et al., 2022)	0.299	0.389	0.422	0.124	0.137	0.145	0.110	0.120	0.126	0.322	0.110
GSCP-Net ($K = 12$)	0.848	**0.860**	**0.861**	0.719	0.832	**0.861**	0.625	0.749	0.854	0.838	0.625
GSCP-Net ($K = 15$)	**0.852**	**0.860**	**0.861**	**0.768**	**0.845**	**0.861**	**0.710**	**0.793**	**0.857**	**0.846**	**0.710**
Source VAE reconstruction	0.861	0.861	0.861	0.861	0.861	0.861	0.861	0.861	0.861	0.861	0.861

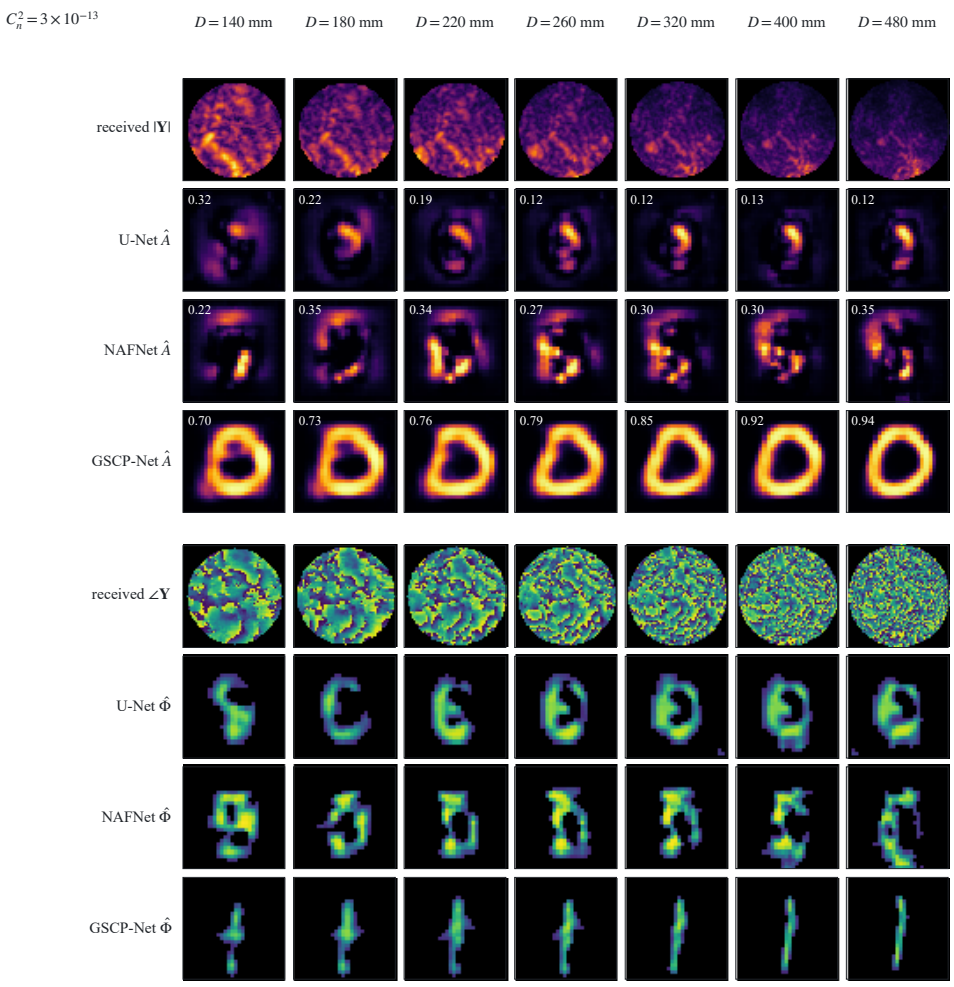


Figure 5: Qualitative aperture sweep at $C_n^2 = 3 \times 10^{-13}\,\mathrm{m}^{-2/3}$. Columns correspond to the seven receiver apertures; rows compare the received field with amplitude and phase reconstructions from U-Net, NAFNet, and GSCP-Net.

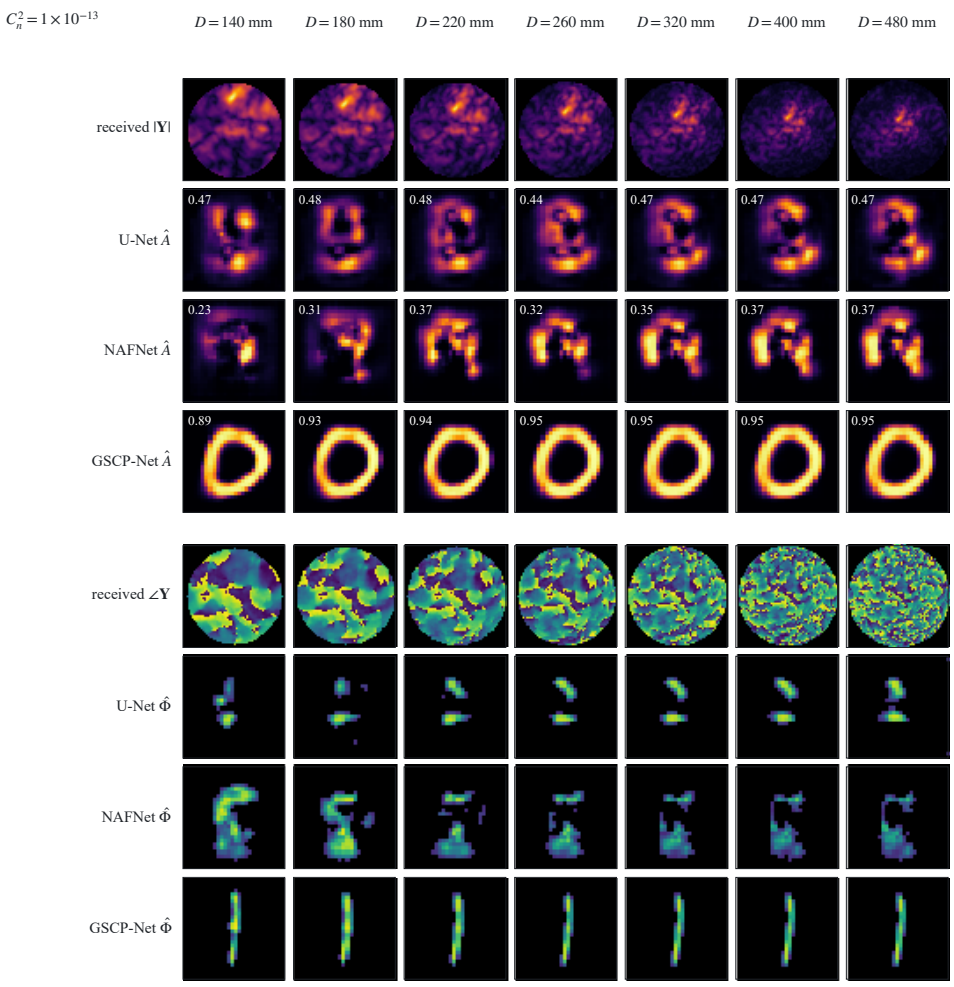


Figure 6: Qualitative aperture sweep at $C_n^2 = 1 \times 10^{-13}\,\mathrm{m}^{-2/3}$. Columns correspond to the seven receiver apertures; rows compare the received field with amplitude and phase reconstructions from U-Net, NAFNet, and GSCP-Net.

C.2 Probe Design and Conditioning

Table 6: Probe-design and probe-number ablation over all 42 turbulence–aperture conditions, providing the complete results corresponding to Table 2.

C_n^2 ($m^{-2/3}$) \ D (mm)	140	180	220	260	320	400	480
Fixed initial probes $Q_k = Q_k^0$ ($K = 12$)							
1×10^{-15}	0.850	0.859	0.861	0.861	0.861	0.861	0.861
3×10^{-15}	0.830	0.853	0.860	0.861	0.861	0.861	0.861
1×10^{-14}	0.773	0.837	0.857	0.860	0.861	0.861	0.861
3×10^{-14}	0.674	0.796	0.841	0.856	0.861	0.861	0.861
1×10^{-13}	0.427	0.619	0.748	0.816	0.850	0.860	0.861
3×10^{-13}	0.302	0.403	0.511	0.603	0.707	0.796	0.837
Learned probes, $K = 9$							
1×10^{-15}	0.846	0.858	0.861	0.861	0.861	0.861	0.861
3×10^{-15}	0.817	0.850	0.860	0.861	0.861	0.861	0.861
1×10^{-14}	0.740	0.829	0.855	0.859	0.861	0.861	0.861
3×10^{-14}	0.585	0.764	0.832	0.852	0.859	0.861	0.861
1×10^{-13}	0.289	0.472	0.650	0.769	0.837	0.858	0.861
3×10^{-13}	0.189	0.259	0.346	0.438	0.568	0.716	0.802

Table 7: Receiver Jacobian conditioning versus probe number K for the fixed initial and learned codebooks. Results report the median and worst-case $\sigma_{\min}(J)$ and $\kappa(J)$ over the 42 turbulence–aperture conditions.

	Fixed initial probes $Q_k = Q_k^0$				Learned probes			
	$\sigma_{\min}(J)$		$\kappa(J)$		$\sigma_{\min}(J)$		$\kappa(J)$	
K	median	worst	median	worst	median	worst	median	worst
9	0.58	0.09	22.8	25.4	0.93	0.15	14.0	14.7
10	1.42	0.22	9.8	10.9	–	–	–	–
11	2.21	0.35	6.6	7.1	–	–	–	–
12	3.06	0.48	5.1	5.5	3.62	0.61	4.2	4.3
13	3.88	0.62	4.2	4.6	–	–	–	–
14	4.65	0.74	3.7	4.0	–	–	–	–
15	5.40	0.86	3.3	3.6	5.53	0.90	3.3	3.3

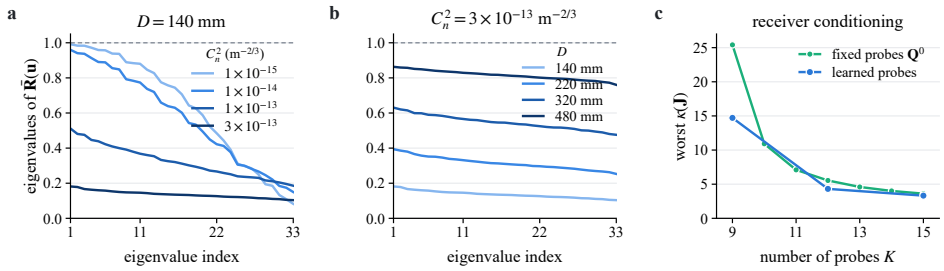


Figure 7: Effective-metric spectra under varying turbulence and aperture conditions, together with the worst-case Jacobian condition number versus probe number for fixed and learned codebooks.

C.3 Dynamic-Turbulence Robustness

Table 8: Dynamic-turbulence recovery fidelity over all tested conditions under Taylor frozen flow, providing the complete results corresponding to Table 3.

C_n^2 ($\mathrm{m}^{-2/3}$), quantity \ D (mm)	140	180	220	260	320	400	480
$C_n^2 = 10^{-14}$							
$t = 0$ (static)	0.839	0.851	0.856	0.857	0.858	0.858	0.858
$t = 2\,t_c$	0.839	0.851	0.856	0.857	0.858	0.858	0.858
$t = 5\,t_c$	0.839	0.852	0.856	0.857	0.858	0.858	0.858
$t = 10\,t_c$	0.840	0.852	0.856	0.857	0.858	0.858	0.858
$t = 20\,t_c$	0.838	0.851	0.855	0.857	0.858	0.858	0.858
$\|R(20t_c) - R(0)\|_F / \|R(0)\|_F$	0.37	0.21	0.10	0.05	0.02	0.00	0.00
$C_n^2 = 3 \times 10^{-14}$							
$t = 0$ (static)	0.798	0.838	0.851	0.855	0.857	0.858	0.858
$t = 2\,t_c$	0.803	0.838	0.850	0.855	0.857	0.858	0.858
$t = 5\,t_c$	0.796	0.836	0.850	0.855	0.857	0.858	0.858
$t = 10\,t_c$	0.792	0.836	0.850	0.855	0.857	0.858	0.858
$t = 20\,t_c$	0.804	0.840	0.852	0.855	0.857	0.858	0.858
$\|R(20t_c) - R(0)\|_F / \|R(0)\|_F$	0.57	0.36	0.20	0.12	0.05	0.01	0.00
$C_n^2 = 10^{-13}$							
$t = 0$ (static)	0.682	0.769	0.819	0.843	0.854	0.857	0.858
$t = 2\,t_c$	0.683	0.770	0.818	0.843	0.854	0.857	0.858
$t = 5\,t_c$	0.689	0.777	0.822	0.843	0.853	0.857	0.858
$t = 10\,t_c$	0.698	0.785	0.825	0.842	0.853	0.857	0.858
$t = 20\,t_c$	0.682	0.780	0.824	0.844	0.854	0.857	0.858
$\|R(20t_c) - R(0)\|_F / \|R(0)\|_F$	0.83	0.61	0.42	0.26	0.12	0.05	0.02
$C_n^2 = 3 \times 10^{-13}$							
$t = 0$ (static)	0.555	0.667	0.745	0.790	0.827	0.848	0.855
$t = 2\,t_c$	0.561	0.675	0.748	0.790	0.827	0.848	0.855
$t = 5\,t_c$	0.561	0.676	0.748	0.792	0.828	0.848	0.855
$t = 10\,t_c$	0.556	0.679	0.751	0.793	0.827	0.847	0.855
$t = 20\,t_c$	0.546	0.672	0.747	0.789	0.824	0.844	0.853
$\|R(20t_c) - R(0)\|_F / \|R(0)\|_F$	0.91	0.75	0.60	0.49	0.34	0.20	0.10

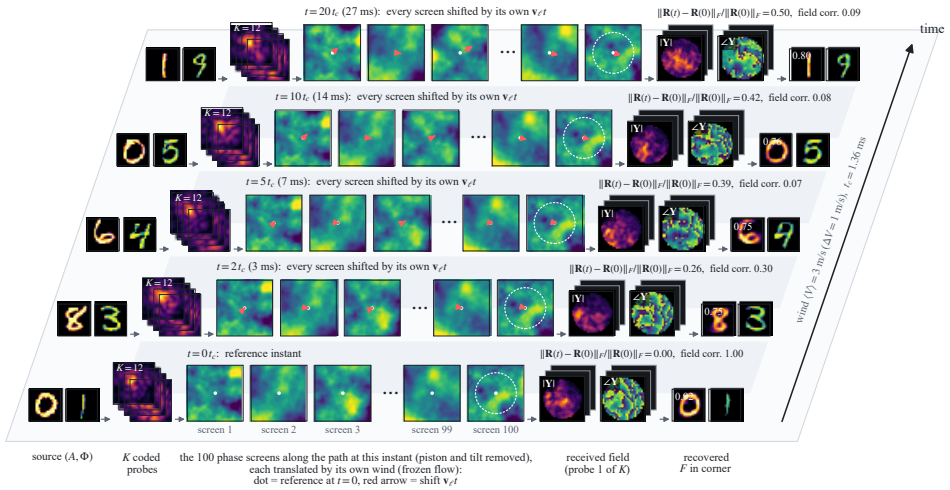


Figure 8: Representative dynamic-turbulence evolution under Taylor frozen flow. The phase screens translate with their individual wind velocities from $t = 0$ to $20t_c$; coded probes, selected phase screens, the received complex field, and the recovered source are shown at each time.

C.4 Generalization and Cross-Dataset Validation

Table 9: Generalization to 30 off-grid turbulence–aperture conditions using interpolated and theoretical mean receiver metrics $\overline{\mathbf{R}}$.

C_n^2 ($\mathrm{m}^{-2/3}$) \ D (mm)	160	200	240	290	360	440
$\overline{\mathbf{R}}$ interpolated from the training grid (used by GSCP-Net)						
2×10^{-15}	0.859	0.861	0.861	0.861	0.861	0.861
6×10^{-15}	0.854	0.859	0.861	0.861	0.861	0.861
2×10^{-14}	0.840	0.855	0.859	0.861	0.861	0.861
6×10^{-14}	0.788	0.830	0.849	0.858	0.861	0.861
2×10^{-13}	0.631	0.723	0.781	0.823	0.848	0.858
Theoretical $\overline{\mathbf{R}}$ at the target condition						
2×10^{-15}	0.859	0.861	0.861	0.861	0.861	0.861
6×10^{-15}	0.854	0.859	0.861	0.861	0.861	0.861
2×10^{-14}	0.841	0.855	0.859	0.861	0.861	0.861
6×10^{-14}	0.790	0.832	0.850	0.858	0.861	0.861
2×10^{-13}	0.623	0.723	0.784	0.825	0.849	0.858

Table 10: Cross-dataset validation of GSCP-Net on EMNIST letters and Fashion-MNIST over the same 42 turbulence–aperture conditions. The corresponding source VAE reconstruction fidelities are 0.871 and 0.848, respectively.

C_n^2 ($\mathrm{m}^{-2/3}$) \ D (mm)	140	180	220	260	320	400	480
EMNIST letters (Cohen et al., 2017)							
1×10^{-15}	0.869	0.870	0.870	0.870	0.871	0.871	0.871
3×10^{-15}	0.865	0.869	0.870	0.870	0.871	0.871	0.871
1×10^{-14}	0.853	0.865	0.869	0.870	0.870	0.871	0.871
3×10^{-14}	0.823	0.855	0.865	0.869	0.870	0.871	0.871
1×10^{-13}	0.679	0.781	0.833	0.855	0.867	0.870	0.870
3×10^{-13}	0.553	0.661	0.734	0.778	0.819	0.849	0.862
Fashion-MNIST (Xiao et al., 2017)							
1×10^{-15}	0.848	0.848	0.848	0.848	0.848	0.848	0.848
3×10^{-15}	0.847	0.848	0.848	0.848	0.848	0.848	0.848
1×10^{-14}	0.843	0.847	0.848	0.848	0.848	0.848	0.848
3×10^{-14}	0.834	0.844	0.847	0.848	0.848	0.848	0.848
1×10^{-13}	0.790	0.824	0.838	0.844	0.847	0.848	0.848
3×10^{-13}	0.741	0.779	0.806	0.821	0.834	0.842	0.846

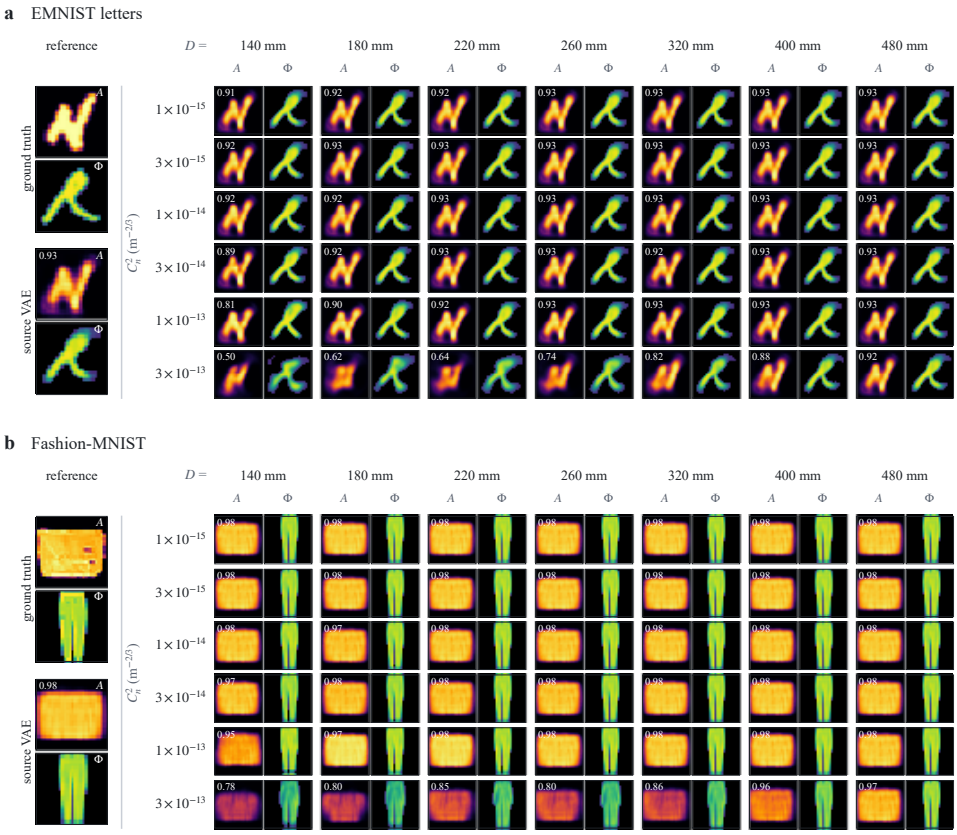


Figure 9: Cross-dataset qualitative recovery on (a) EMNIST letters and (b) Fashion-MNIST over the same turbulence–aperture grid. Reference and source-VAE fields are shown together with the recovered amplitude A and phase Φ.

C.5 Physical Conditions and Implementation Details

Table 11: Physical parameters of the 42 turbulence–aperture conditions for $L = 2$ km and $\lambda = 1550$ nm, including r_0, σ_R^2, t_c for $\langle V \rangle = 3$ m/s, and D/r_0.

				D/r_0 for D (mm) =						
C_n^2 (m$^{-2/3}$)	r_0 (mm)	σ_R^2	t_c (ms)	140	180	220	260	320	400	480
1×10^{-15}	206.1	0.07	21.60	0.68	0.87	1.07	1.26	1.55	1.94	2.33
3×10^{-15}	106.6	0.21	11.17	1.31	1.69	2.06	2.44	3.00	3.75	4.50
1×10^{-14}	51.8	0.71	5.43	2.70	3.48	4.25	5.02	6.18	7.73	9.27
3×10^{-14}	26.8	2.13	2.81	5.23	6.72	8.21	9.71	11.9	14.9	17.9
1×10^{-13}	13.0	7.09	1.36	10.8	13.8	16.9	20.0	24.6	30.8	36.9
3×10^{-13}	6.7	21.28	0.71	20.8	26.8	32.7	38.6	47.6	59.5	71.3

Table 12: Implementation summary of GSCP-Net and the end-to-end baselines. All methods use the same propagation model, turbulence–aperture condition grid, and paired test protocol. Training times are measured on a single RTX 2060 GPU. For GSCP-Net, only the probe codebook depends on the channel; the channel-independent source VAE (5.45×10^6 parameters) is trained once beforehand, and its ≈5-min training is included in the reported time.

	GSCP-Net ($K = 12$)	U-Net (Tan et al., 2022)	NAFNet (Chen et al., 2022)
Transmitted field	coded probes $\mathbf{B}\mathbf{Q}_k\mathbf{a}(\mathbf{z})$	source field $\mathbf{U}_0$	as U-Net
Receiver input	Gram features of received coded probes	received complex field	as U-Net
Output	latent code $\widehat{\mathbf{z}}$ decoded by the frozen VAE	amplitude–phase maps	amplitude–phase maps
Reconstruction	physics-based latent inversion with trainable unitary probes	end-to-end U-Net	end-to-end NAFNet
Channel-dependent trainable parameters	1.2×10^4	7.77×10^6	4.01×10^6
Training time	≈40 min	≈ 2.9 h	≈4.0 h
Evaluation protocol	Same 42 turbulence–aperture conditions and paired test benchmark		