
Unlocking Inverse Problems Using Deep Learning: Breaking Symmetries in Phase Retrieval

Kshitij Tayal¹, Chieh-Hsin Lai², Raunak Manekar¹, Zhong Zhuang³, Vipin Kumar¹, Ju Sun¹

¹Department of Computer Science and Engineering, University of Minnesota, Twin Cities, USA

²School of Mathematics, University of Minnesota, Twin Cities, USA

³Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities, USA

Abstract

In many physical systems, inputs related by intrinsic system symmetries generate the same output. So when inverting such systems, an input is mapped to multiple symmetry-related outputs. This causes fundamental difficulty for tackling these inverse problems by the emerging end-to-end deep learning approach. Taking phase retrieval as an illustrative example, we show that careful symmetry breaking on the training data can help get rid of the difficulty and significantly improve learning performance on real data.

1 Introduction

For many physical systems, we observe only the output and strive to infer the input. The inference task is called *inverse problem*. Formally, the underlying system is modeled by a forward mapping f , and solving the inverse problem amounts to identifying the inverse mapping f^{-1} .

Let \mathbf{y} denote the observed output. Traditionally, inverse problems are phrased as regularized optimization problems: $\min_{\mathbf{x}} \ell(\mathbf{y}, f(\mathbf{x})) + \lambda\Omega(\mathbf{x})$, where \mathbf{x} represents the input to be estimated, $\ell(\mathbf{y}, f(\mathbf{x}))$ ensures $\mathbf{y} \approx f(\mathbf{x})$, and $\Omega(\mathbf{x})$ encodes structural information about \mathbf{x} to make the problem well posed. Deep learning has enabled the use of data-driven ℓ or Ω , or replacing mappings in iterative methods for solving the regularized formulation by data-adaptive ones. The most radical is using neural networks directly to approximate f^{-1} , also known as the *end-to-end approach*. Several review articles [1–4] have covered these recent developments. We note that most of the successes are about *linear* inverse problems, i.e., f is linear.

In this paper, we focus on the end-to-end approach applied to non-linear inverse problems, and take phase retrieval (PR)—which is central to scientific imaging [5]—as an example. Here, we focus on the 2D version, which is seen most frequently in applications. Given $\mathbf{X} \in \mathbb{C}^{n_1 \times n_2}$, the forward mapping is

$$\mathbf{Y} = \left| \mathbf{F}_{m_1 \times n_1} \mathbf{X} \mathbf{F}_{m_2 \times n_2}^\top \right|^2 \in \mathbb{R}^{m_1 \times m_2}, \quad (1.1)$$

where $\mathbf{F}_{m_1 \times n_1}$ consists of the first n_1 columns of the Fourier matrix $\mathbf{F}_{m_1 \times m_1}$, often called partial Fourier matrix; similarly for $\mathbf{F}_{m_2 \times n_2}$. The mapping is *generically* injective when $m_1 \geq 2n_1 - 1$ and $m_2 \geq 2n_2 - 1$, up to three intrinsic symmetries: 1) 2D translation of the nonzero content of \mathbf{X} ; 2) conjugate of 2D flipping of \mathbf{X} ; and 3) global phase transfer to \mathbf{X} : $\mathbf{X}e^{i\theta}$ for any $\theta \in [-\pi, \pi)$. Any composition of these changes to \mathbf{X} will leave the observation \mathbf{Y} unaltered. Fig. 1 illustrate the first two symmetries, assuming \mathbf{X} is a real-valued image.

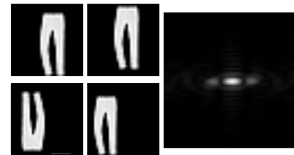


Figure 1: Symmetries in 2D PR. (Left) shifted and flipped copies of the same image; (Right) their common Fourier magnitude

Symmetries can cause significant difficulty for the end-to-end approach. To see this, suppose we randomly sample real values x_i 's and form a training set $\{x_i, x_i^2\}$ and try to learn the square-root function, allowing both positive and negative outputs, using the end-to-end approach. Now if we think of the function determined by the training set, which the neural network is trying to approximate, it is highly oscillatory (see Fig. 2): the sign symmetry in the forward mapping $x_i \mapsto x_i^2$ dictates that in the training set, there are frequent cases where x_i^2 and x_j^2 are close but x_i and x_j have different signs and are far apart. Although in theory neural networks with adequate capacity are universal function approximators, in practice they will struggle to learn such irregular functions. For general inverse problems, so long as the forward symmetries can relate remote inputs to the same output, similar problems can surface. Solving PR is not free from the trouble.

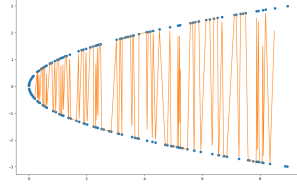


Figure 2: Learn to take square root. An oscillatory function (in orange) determined by the training set. Image credit: [6]

2 Breaking symmetries for PR

The above difficulty has recently been articulated in [6, 7]. Particularly, [6] has proposed a method to resolve the difficulty in Gaussian PR—a simplified version of PR—called *symmetry breaking*. The proposed working principle is to identify a connected (in topological sense), smallest representative subset of the input space, so that all training samples are mapped to such subset to induce a smoother function to be approximated. Take the above square root example again, the positive ray $\mathbb{R}_+ \subset \mathbb{R}$ can be taken, as it is topologically connected, and any $x \in \mathbb{R}$ (except 0) can be represented by an element in \mathbb{R}_+ by a sign flipping, and it cannot be made smaller to remain representative. Then the training data set $\{x_i, x_i^2\}$ will be processed so that any pair (x_i, x_i^2) with negative x_i will be changed to $(-x_i, x_i^2)$, with the rest data points unchanged. Obviously after this processing to the training set, the function determined will closely trace the upper branch of the square root function, which is much more smooth.

PR has three symmetries as discussed above. Under the global phase transfer, equivalent data points form continuous curves that are easy to represent algebraically. The conjugate 2D flipping and nonzero content translation, however, induce irregular equivalent sets that are hard to represent. Following [6] and prescribing a rule for symmetry breaking in the original \mathbf{X} space seems hopeless.

Fortunately, the three symmetries can be equivalently represented in the complex phase $e^{i\theta}$ space after the Fourier transform. Let \mathcal{X} denote the oversampled Fourier transform of \mathbf{X} . Now 1) for 2D translation, any allowable 2D translation $t_1, t_2 \in \mathbb{Z}$ induces the change $\mathcal{X}(k_1, k_2) \mapsto e^{i2\pi(\frac{k_1 t_1}{m_1} + \frac{k_2 t_2}{m_2})} \mathcal{X}(k_1, k_2)$; 2) conjugate 2D flipping induces the change $\mathcal{X} \mapsto \overline{\mathcal{X}}$, i.e., change to the complex phase $e^{i\Theta} \mapsto e^{-i\Theta}$; and 3) global phase transfer induces the change $\mathcal{X} \mapsto e^{i\theta} \mathcal{X}$. The change due to 2) is a global sign flipping in the *angle* space, and the equivalent set due to 3) is a line in the *angle* space. But 1) is still relatively irregular whether represented in the angle or phase space.

Our strategy here is a combination of *rigorous* symmetry breaking for 2) and 3) in the complex phase space and *heuristic* symmetry breaking for 1) in the original space—our later real-data experiments confirm that the combination is effective. To break 1), we propose to simply center the nonzero content as a heuristic. To break 2) and 3), we perform a geometric construction of a connected, smallest representative subset in the angle space and then represent it in the phase space to avoid the tricky 2π periodicity issue in the angle space. For brevity, we omit intuition behind the construction and directly present the results as follows.

Consider the following set in the phase domain

$$\mathcal{R} \doteq \left\{ \Omega \in \mathbb{C}^{m_1 \times m_2} : \Omega(1, 1) = 1, \Omega(1, 2) \in \mathbb{S}_+, \Omega(i, j) \in \mathbb{S} \forall \text{ other index } (i, j) \right\}, \quad (2.1)$$

where \mathbb{S} denotes the 1D complex circle and \mathbb{S}_+ the upper half circle. We can prove the following, stated in the equivalent vector space for convenience. We write $\mathcal{R} \subset \mathbb{S}^{m_1 m_2}$ to mean the equivalent of \mathcal{R} in the vector notation.

Theorem 2.1. *Consider the conjugate flipping and global phase transfer symmetries only. The set \mathcal{R} is a connected, smallest representative in the phase domain $\mathbb{S}^{m_1 m_2}$ with a negligible set $\mathcal{N} = \{1\} \times \{\omega \in \mathbb{S} : \text{Im}(\omega) = 0\}^{m_1 m_2 - 1}$.*

Proof. See Appendix A.1. □

To apply this, we work with end-to-end DNNs that directly predicts the $m_1 \times m_2$ complex phases. We first center the nonzero content inside \mathbf{X}_i 's in the training set, and then take the oversampled Fourier transform and perform the symmetry breaking as implied by Theorem 2.1 in the complex phase space. For any phase matrix Ω , the symmetry breaking goes naturally as follows: first a global phase transfer is performed to make $\Omega(1, 1) = 1$, and then a global angle (here we assume the angle has been transferred to the range of $[-\pi, \pi)$) negation is performed, i.e., $\Omega \mapsto -\Omega$ if the second angle is negative.

Symmetry breaking for general inverse problems For general inverse problems, although the symmetries might be very different than here and the sample spaces could also be more complicated, the three properties, which concern only the geometric and topological aspects of the space, can be generalized as a basic mathematical principle for effective symmetry breaking. Our symmetry-breaking solution for PR also suggests that for problems with multiple symmetries, one may need to look at a transformed space, or even mixture of spaces for different symmetries for efficient representation and symmetry breaking.

3 Numerical Experiments

In this section, we set up a preliminary experiment to verify our claim that effective symmetry breaking facilitates efficient learning. Particularly, we show that symmetry breaking substantially improves PR performance over alternative methods.

We conduct our experiments on the Fashion-MNIST dataset [8]. We take their 60,000 training images and 10,000 test images to construct our training and test sets respectively. Each example is a 28×28 grayscale image. To simulate the typical black ground that causes the translation freedom in PR applications, we place all the images in a black background of 42×42 —most previous methods overlook this in their experiments, but practically the translation freedom, or what PR community call support estimation, is a major failing factor for most PR methods. So $n = 42$, and we take $m = 96$ here to ensure injectivity of the forward model $2n - 1 = 83$ is exceeded. We create 4 variants of the dataset to test the impact of symmetries on learning—this is the first time this kind of rigorous evaluation is performed; most previous methods use natural image datasets where the image contents are naturally centered and oriented, which does not match the scenarios in PR applications e.g., in coherent diffraction imaging. We do this by modifying the images as described below, followed by the standard operation of taking Fourier magnitudes.

- **No Symmetry:** i.e., all images are placed in the center of the black background i.e. padding 7 pixels on all side; samples shown in Fig. 3 (a)-left;
- **Flipping symmetry:** all images are placed in the center of the black background and 50% of randomly selected training and test images are 2D flipped; samples shown in Fig. 3 (b)-left.
- **Shift symmetry:** all images placed in a larger dark background and randomly translated; samples shown in Fig. 3 (c)-left;
- **Shift and Flipping symmetries:** random flipping followed by random translation; samples shown in Fig. 3 (d)-left.

Results on randomly selected test images are presented in Fig. 3. We use U-Net [9] as our backbone neural network. For each variant of the dataset, the left column is the groundtruth image, and the middle and right columns are results produced by U-Net-*B*, i.e., without symmetry breaking—this is exactly the method used in [10], one of the state-of-the-art methods based on the end-to-end approach, and U-Net-*A*, i.e., with symmetry breaking—our method, respectively.

Table 1: Test error (MSE) using different symmetry schemes

	U-Net- <i>B</i>	U-Net- <i>A</i> (ours)
No Symmetry	0.103	0.103
Flipping Symmetry	0.168	0.162
Shift Symmetry	0.249	0.102
Shift & Flipping Symmetry	0.248	0.161

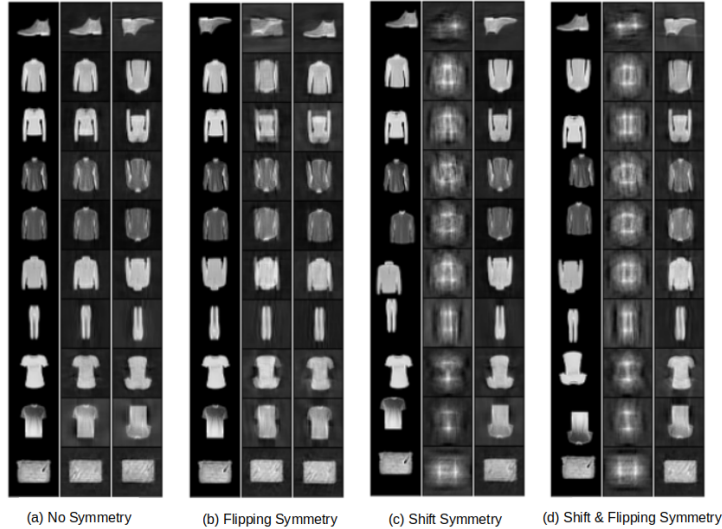


Figure 3: Visualization of recovery results. For each group, the 1st column contains the groundtruth images, and the 2nd and 3rd columns are reconstructions produced by U-Net-B and U-Net-A, respectively. While U-Net-B fails to perform right reconstruction once symmetries are built into the dataset, our U-Net-A produces consistent recovery results with or without the symmetric data.

First note that when no explicit symmetries are built into the dataset, U-Net-B, a representative end-to-end method for PR [10], gives good recovery. But it fails once we build in the essential symmetries. The mode of failure is interesting, as the estimated images are almost always the superposition of the symmetric (translated or flipped) copies of the groundtruth. This is very similar to the failure mode of the classic iterative methods on PR. Moreover, for images that are visually similar between the original and the flipped copy e.g: “handbag”, “leggings”, the reconstruction results are good with or without the flipping symmetry, consistent with our intuition.

Table 2: MSE error

Method	MSE
ALM	0.299
U-Net-B	0.249
U-Net-A	0.160

On the other hand, irrespective of the symmetries, U-Net-A consistently leads to good recovery. Table 1 provides the average MSE adjusted to the symmetries (defined in Appendix A.2) for the test set. As noted above, absent symmetries, both U-Net-B and U-Net-A work well and the average MSEs are the same. However, once the dataset contains the essential symmetries, we see a substantial gap in the MSEs of the reconstructed images, which is consistent with the visual results.



Figure 4: Comparison with ALM. 1st column contains groundtruth images, 2nd, 3rd and 4th column are reconstructions from ALM, U-Net-B and U-Net-A respectively.

For practical PR, mostly iterative methods are deployed. However, these methods are known to fail when there is translation freedom in the image and the support (i.e., location of nonzero pixels) of the image content is not precisely known. To see if our end-to-end approach makes progress on this, we compare it with a state-of-the-art iterative method for PR recently proposed in [11] that has demonstrated good numerical stability and competitive performance, dubbed ALM. Here we only experiment with the most realistic version of the dataset, i.e., with both shift and flipping symmetries. Results on randomly selected test images are presented in Fig. 4 and quantitatively results are presented in Table 2. Visually, our method faithfully reconstructs the holistic content of the original images, whereas both U-Net-B and ALM fail miserably. Quantitatively, our method leads the other two by a considerable gap in MSE.

These results show that symmetry breaking is significant in unlocking the true potential of the end-to-end approach for solving PR in particular, and nonlinear inverse problems with symmetries in general. In a companion paper [12], we describe a less mathematical yet effective way of performing

symmetry breaking for PR, that also seems to be generally applicable for other nonlinear inverse problems.

Acknowledgments and Disclosure of Funding

KT and VK are supported by NSF HDR 1934721. The authors acknowledge the Minnesota Supercomputing Institute (MSI) at the University of Minnesota for providing resources that contributed to the research results reported within this paper.

References

- [1] M. T. McCann, K. H. Jin, and M. Unser, “Convolutional neural networks for inverse problems in imaging: A review,” *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 85–95, nov 2017.
- [2] A. Lucas, M. Iliadis, R. Molina, and A. K. Katsaggelos, “Using deep neural networks for inverse problems in imaging: Beyond analytical methods,” *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 20–36, jan 2018.
- [3] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb, “Solving inverse problems using data-driven models,” *Acta Numerica*, vol. 28, pp. 1–174, may 2019.
- [4] G. Ongie, A. Jalal, C. A. Metzler, R. G. Baraniuk, A. G. Dimakis, and R. Willett, “Deep learning techniques for inverse problems in imaging,” *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 1, pp. 39–56, may 2020.
- [5] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, “Phase retrieval with application to optical imaging: a contemporary overview,” *IEEE signal processing magazine*, vol. 32, no. 3, pp. 87–109, 2015.
- [6] K. Tayal, C.-H. Lai, V. Kumar, and J. Sun, “Inverse problems, deep learning, and symmetry breaking,” *arXiv preprint arXiv:2003.09077*, 2020.
- [7] R. Manekar, K. Tayal, V. Kumar, and J. Sun, “End-to-end learning for phase retrieval,” *ICML workshop on ML Interpretability for Scientific Discovery*, 2020.
- [8] H. Xiao, K. Rasul, and R. Vollgraf. (2017) Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms.
- [9] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [10] A. Sinha, J. Lee, S. Li, and G. Barbastathis, “Lensless computational imaging through deep learning,” *Optica*, vol. 4, no. 9, p. 1117, sep 2017.
- [11] Z. Zhuang, G. Wang, Y. Travadi, and J. Sun, “Phase retrieval via second-order nonsmooth optimization,” *ICML workshop on Beyond First Order Methods in Machine Learning*, 2020.
- [12] R. Manekar, Z. Zhuang, K. Tayal, V. Kumar, and J. Sun, “Deep learning initialized phase retrieval,” *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems*, 2020.

A Deferred Proofs

A.1 Proof of Theorem 2.1

Proof. It is clear that \mathcal{R} is connected in $\mathbb{S}^{m_1 m_2}$ since \mathcal{R} is path-connected set on $\mathbb{S}^{m_1 \times m_2}$ with the inherited subspace Euclidean topology of $\mathbb{C}^{m_1 \times m_2}$. Also, \mathcal{N} is of Lebesgue measure 0 since it is a product of finite points. Now we are going to prove \mathcal{R} is a representative of $\mathbb{S}^{m_1 m_2}$.

Let \mathfrak{G} be the set of all possible symmetry transfers composed of sequences of global phase transfer and global phase conjugation in PR. For any given $\mathbf{z} = (e^{i\theta_0}, e^{i\theta_1}, \dots, e^{i\theta_{m_1 m_2 - 1}})$, we need to find a $\omega \in \mathcal{R}$ such that there is a $g \in \mathfrak{G}$ satisfying $g(\omega) = \mathbf{z}$. If $\text{Im}(e^{i(\theta_1 - \theta_0)}) > 0$, we take $\omega = (1, e^{i(\theta_1 - \theta_0)}, e^{i(\theta_2 - \theta_0)}, \dots, e^{i(\theta_{m_1 m_2 - 1} - \theta_0)})$ then $\omega \in \mathcal{R}$ and $e^{i\theta_0} \omega = \mathbf{z}$. On the other hand, if $\text{Im}(e^{i(\theta_1 - \theta_0)}) < 0$, we can consider the conjugate format $\omega =$

$(1, \overline{e^{i(\theta_1 - \theta_0)}}, \overline{e^{i(\theta_2 - \theta_0)}}, \dots, \overline{e^{i(\theta_{m_1 m_2 - 1} - \theta_0)}}) \in \mathcal{R}$ and obviously a global phase negation followed by a global phase transfer $e^{i\theta_0}$ leads to \mathbf{z} . This proves that \mathcal{R} is representative.

At last, we need to show the smallestness in the sense that with any point of \mathcal{R} removed, we cannot recover it by other points in \mathcal{R} . That is, with arbitrary $\tilde{\mathbf{z}} \in \mathcal{R}$ given, for all $g \in \mathfrak{G}$ and all $\mathbf{z} \in \mathcal{R} \setminus \{\tilde{\mathbf{z}}\}$, we have $g(\mathbf{z}) \neq \tilde{\mathbf{z}}$.

We first claim that any g is equivalent to an optional global phase conjugation followed by a global phase transfer. To see this, it is sufficient to prove that the order of phase conjugation and phase transfer can be exchanged. Let ψ denote a global phase transfer by $e^{i\psi}$ and f phase conjugation. Now if $\psi \circ f = f \circ \psi'$, or $-(\psi' + \theta) = -\theta + \psi + 2k\pi$, we have $\psi' = -\psi - 2k\pi$. So one can keep exchanging conjugation and transfer so that all conjugations precede transfers. The conjugations now can be equivalently written as an optional conjugation, and the transfers as a single transfer.

Now we can go back to the proof of smallestness. Write $\tilde{\mathbf{z}} = (e^{i\tilde{\theta}_0}, e^{i\tilde{\theta}_1}, \dots, e^{i\tilde{\theta}_{m_1 m_2 - 1}})$ and $\mathbf{z} = (e^{i\theta_0}, e^{i\theta_1}, \dots, e^{i\theta_{m_1 m_2 - 1}})$ where $\tilde{\theta}_0 = \theta_0 = 0$ and $\text{Im}(e^{i\tilde{\theta}_1}), \text{Im}(e^{i\theta_1}) > 0$. Suppose that there is a $g \in \mathfrak{G}$ such that $\tilde{\mathbf{z}} = g(\mathbf{z})$. we may assume $g = f \circ \psi$ or $g = \psi$ where ψ is a phase transition with the total angles ψ and f is the conjugate flipping. If $g = f \circ \psi$, $\tilde{\mathbf{z}} = g(\mathbf{z})$ implies that

$$\tilde{\theta}_j \equiv -(\psi + \theta_j) + 2\pi k_j \pmod{2\pi} \quad \forall j \quad (\text{A.1})$$

for some $k_j \in \mathbb{Z}$. We can solve $\psi = 2\pi k_0$ as $j = 0$ and this implies $\tilde{\theta}_j \equiv 2\pi(k_j - k_0) - \theta_j \pmod{2\pi} \equiv -\theta_j$ for all j , especially, $\tilde{\theta}_1 = -\theta_1$. This contradict with the fact that $\text{Im}(e^{i\tilde{\theta}_1}), \text{Im}(e^{i\theta_1}) > 0$. If $g = \psi$, we then have the relationship

$$\tilde{\theta}_j \equiv (\psi + \theta_j) + 2\pi k_j \pmod{2\pi}. \quad (\text{A.2})$$

Again, we can solve $\psi = -2\pi k_0$ as $j = 0$ and this indicates that $\tilde{\mathbf{z}} = \mathbf{z}$ which contradicts the assumption. Hence, we prove the smallestness. \square

A.2 Mean Square Error (MSE)

Our reconstructed image is in $\mathbb{C}^{m \times m}$, where our original image is in $\mathbb{C}^{n \times n}$. To account for the three symmetries when taking MSE measure, we take the following steps: we take the original image, and scan through the larger reconstructed image to account for the translation symmetry. At each scan position, we calculate an adjusted MSE between the current patch $\mathbf{B} \in \mathbb{C}^{n \times n}$ and the original image \mathbf{A} . A $\lambda > 0$ and a global phase factor $e^{i\theta}$ (to account for the global phase) are introduced when calculating the MSE, i.e.,

$$\min_{\theta, \eta \geq 0} \|\mathbf{A} - \eta \mathbf{B} e^{i\theta}\|_F^2. \quad (\text{A.3})$$

The smallest adjusted MSE is recorded over all scan positions. Then, the original image \mathbf{A} is 2D flipped and the same scanning process is repeated to calculate another smallest MSE, to account for the flipping symmetry. The smaller of the smallest MSE values is finally taken.

Below, we show that the optimal value in Eq. (A.3) can be easily computed. First we expand the square inside the objective and perform partial minimization with respect to θ , leading to

$$\max_{\theta} \text{Re} \langle \mathbf{A}, \mathbf{B} e^{i\theta} \rangle. \quad (\text{A.4})$$

But $\text{Re} \langle \mathbf{A}, \mathbf{B} e^{i\theta} \rangle = \text{Re} (\langle \mathbf{A}, \mathbf{B} \rangle e^{i\theta}) \leq |\langle \mathbf{A}, \mathbf{B} \rangle e^{i\theta}| \leq |\langle \mathbf{A}, \mathbf{B} \rangle|$ and the upper bound is achievable when $\theta = -\angle \langle \mathbf{A}, \mathbf{B} \rangle$. So the optimization problem now becomes

$$\min_{\eta \geq 0} \|\mathbf{A}\|_F^2 + \eta^2 \|\mathbf{B}\|_F^2 - 2\eta |\langle \mathbf{A}, \mathbf{B} \rangle|. \quad (\text{A.5})$$

The minimum of Eq. (A.5) occurs either when $\eta = 0$, which is $\|\mathbf{A}\|_F^2$, or when $2\eta \|\mathbf{B}\|_F^2 = 2|\langle \mathbf{A}, \mathbf{B} \rangle| \implies \eta = |\langle \mathbf{A}, \mathbf{B} \rangle| / \|\mathbf{B}\|_F^2$, leading to the function value

$$\|\mathbf{A}\|_F^2 - \frac{|\langle \mathbf{A}, \mathbf{B} \rangle|^2}{\|\mathbf{B}\|_F^2}, \quad (\text{A.6})$$

which is the smaller one.