Phase Retrieval Using Conditional Generative Adversarial Networks

Tobias Uelwer, Alexander Oberstraß, Stefan Harmeling

25th INTERNATIONAL CONFERENCE ON PATTERN RECOGNITION
Milan, Italy 10 | 15 January 2021

Heinrich Heine Universität Düsseldorf

Fourier Phase Retrieval

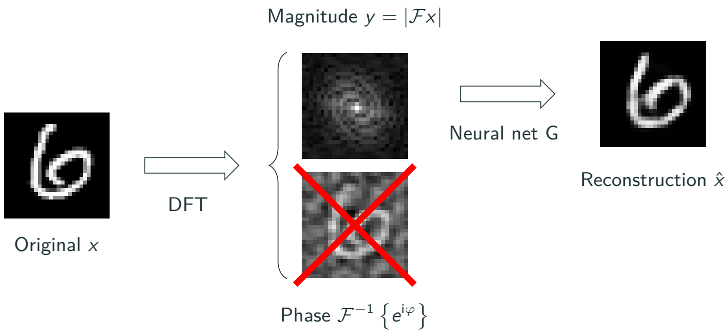

- Phase retrieval aims at recovering a signal x from its Fourier magnitudes $y = |\mathcal{F}x|$, where $\mathcal{F}$ is the Fourier transform.
- Fourier phase retrieval arises in different areas of research, for example in X-ray crystallography, microscopy or astronomical imaging.
- This problem is ill-posed since many signals can have the same Fourier magnitudes.

Proposed Method: Conditional GANs for Phase Retrieval

We train a conditional GAN [3, 5] to reconstruct the images from their Fourier magnitude.

- **At training time** the following optimization problem is solved:

$$\min_G \max_D \mathcal{L}_{\text{adv}}(D, G) + \lambda \mathcal{L}_{\text{rec}}(G), \tag{1}$$

with an adversarial component:

$$\mathcal{L}_{\text{adv}}(D, G) = \mathbb{E}_x \big[\log D(x, y) \big] + \mathbb{E}_{x,z} \big[\log \big(1 - D(G(z, y), y) \big) \big], \tag{2}$$

and a reconstruction component

$$\mathcal{L}_{\text{rec}}(G) = \mathbb{E}_{x,z} \big[\| x - G(z, y) \|_1 \big]. \tag{3}$$

Proposed Method: Conditional GANs for Phase Retrieval

- **At test time** the latent variable z is optimized for each measurement y to minimize the error

$$z^* = \arg\min_z \left\| y - |\mathcal{F} G(z, y)| \right\|_2^2 . \tag{4}$$

  to find an $\hat{x} = G(z^*, y) \approx x$. This is inspired by the work of Hand et al. [2].
- We denote our approach with this latent-space optimization as PRCGAN*, whereas the approach without this optimization as PRCGAN.

Evaluation: Experimental Setup

For the evaluation we consider the following experimental setup:

- We compare the performance with the following existing methods:
  1. Fienup's hybrid-input-output (HIO) algorithm [1]
  2. Relaxed Averaged Alternating Reflections (RAAR) algorithm [4]
  3. End-to-end (E2E) learning with MAE-loss function
  4. Deep Generative Prior (DPR) [2].
- We use the following datsets: MINIST, Fashion-MNIST and CelebA.
- We report the mean squared error (MSE), the mean absolute error (MAE) and the structural similarity index (SSIM) of the reconstructions.

Evaluation: Fourier PR on MNIST

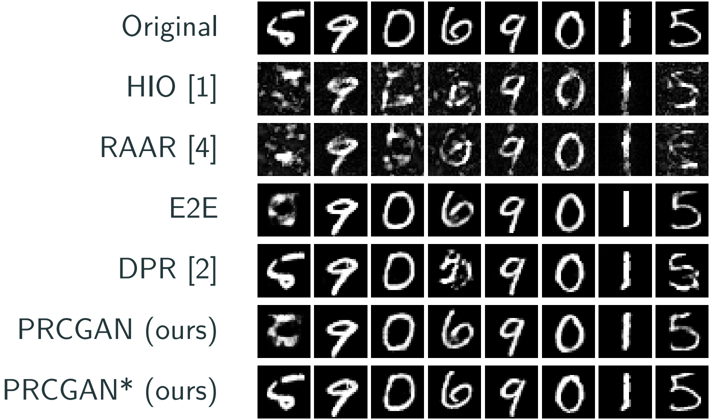

Figure 1: Reconstructions from the Fourier magnitudes of the MNIST test dataset.

Evaluation: Fourier PR on Fashion-MNIST

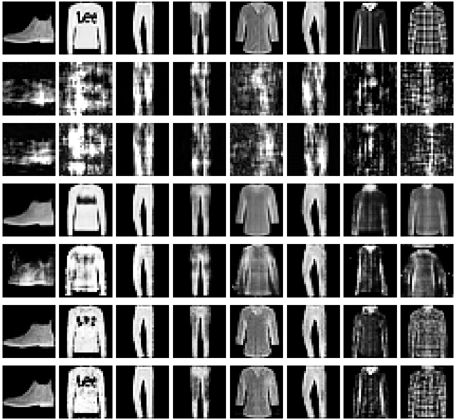

Figure 2: Reconstructions from the Fourier magnitudes of the Fashion-MNIST test dataset.

Evaluation: Fourier PR on MNIST and Fashion-MNIST

Dataset	Metric	HIO	RAAR	E2E	DPR	PRCGAN	PRCGAN*
MNIST	MSE	0.0441	0.0489	0.0183	0.0093	0.0168	**0.0010**
	MAE	0.1016	0.1150	0.0411	0.0221	0.0399	**0.0043**
	SSIM	0.5708	0.5232	0.8345	0.9188	0.8449	**0.9898**
Fashion-MNIST	MSE	0.0646	0.0669	0.0128	0.0280	0.0151	**0.0087**
	MAE	0.1604	0.1673	0.0526	0.0856	0.0572	**0.0412**
	SSIM	0.4404	0.4314	0.7940	0.6602	0.7749	**0.8580**

Table 1: Evaluation results for MNIST, Fashion-MNIST and CelebA for the reconstructions from the Fourier magnitudes. We register the reconstructions for MNIST and Fashion-MNIST.

Evaluation: Fourier PR CelebA

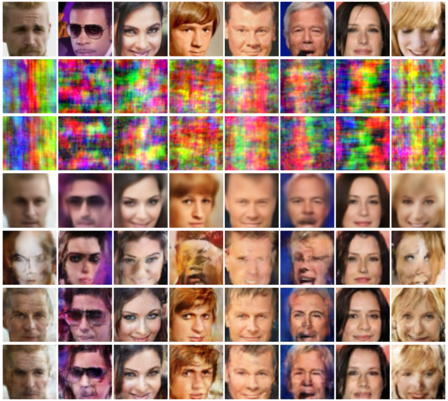

Figure 3: Reconstructions from Fourier magnitudes of the CelebA test dataset.

Evaluation: Fourier Phase Retrieval CelebA

Dataset	Metric	HIO	RAAR	E2E	DPR	PRCGAN	PRCGAN*
	MSE	0.0737	0.0729	0.0106	0.0388	0.0138	**0.0093**
CelebA	MAE	0.2088	0.2073	0.0699	0.1323	0.0804	**0.0642**
	SSIM	0.1671	0.2274	0.7444	0.5299	0.6799	**0.7631**

Table 2: Evaluation results for MNIST, Fashion-MNIST and CelebA for the reconstructions from the Fourier magnitudes. We register the reconstructions for MNIST and Fashion-MNIST.

References

James R Fienup.
Phase retrieval algorithms: a comparison.
Applied optics, 21(15):2758–2769, 1982.

Paul Hand, Oscar Leong, and Vlad Voroninski.
Phase retrieval under a generative prior.
pages 9136–9146, 2018.

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros.
Image-to-image translation with conditional adversarial networks.
In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.

D Russell Luke.
Relaxed averaged alternating reflections for diffraction imaging.
Inverse problems, 21(1):37–50, 2005.

Mehdi Mirza and Simon Osindero.
Conditional generative adversarial nets.
arXiv preprint arXiv:1411.1784, 2014.