Robust Phase Retrieval with Prior Knowledge

Anqi Fu
anqif@stanford.edu

October 26, 2016

1 Introduction

In many imaging applications, one can only measure the power spectral density (the magnitude of the Fourier transform) and not the phase of the signal of interest. For example, in optical settings, devices like CCD cameras only capture the photon flux, so important information stored in the phase is lost. The challenge is to reconstruct the original signal from measurements of its Fourier magnitude. This problem, known as phase retrieval, arises in many scientific and engineering applications, including optics, crystallography, and astronomy.

2 Problem Description

Suppose $x = (x_1, \ldots, x_N)$ is a non-negative signal, and denote by $\mathcal{F} : x \rightarrow F$ the operator that returns its N-point discrete Fourier transform (DFT). We observe $b = (b_1, \ldots, b_N) \in \mathbb{R}_+^N$, the magnitude of the Fourier transform of the signal, and wish to find x that satisfies $|\mathcal{F}x| = b$. Phase retrieval can be posed as the optimization problem

$$\min_x \{f(|\mathcal{F}x| - b)\} \tag{1}$$

3 Project Goals

My goal is to reconstruct the signal $x \in \mathbb{R}_+^N$ as accurately as possible, where accuracy is determined by the peak signal-to-noise ratio (PSNR). First, I will implement the classic error reduction (ER) and hybrid input-output (HIO) algorithms described in [SWL16]. At every iteration k, these algorithms solve 1 with $f(x) = ||x||_2^2$ via a gradient descent update followed by a projection onto the support, i.e. S^k such that $x_i = 0$ for all $i \notin S^k$. For pixels outside the support, ER sets $x_i^{k+1} = 0$ uniformly, while HIO incorporates feedback information from the current iteration's x_i^k.

Then, time permitting, I will consider two extensions of the optimization problem: a weighted prior and a regularization function with a known proximal gradient. In the first case, I include a total variation penalty on x

$$\min_x \left\{ \alpha f(|\mathcal{F}x| - b) + (1-\alpha) \sum_{n=1}^{N-1} (x_{n+1} - x_n) \right\}, \quad s.t. \;\; x \geq 0 \tag{2}$$

where the weight $\alpha \in [0, 1]$ is chosen by cross-validation. In the second case, I add a regularization term $r(x)$

$$\min_x \{ f(|\mathcal{F}x| - b) + r(x) \}, \quad s.t. \;\; x \geq 0 \tag{3}$$

where the proximal operator of r is defined as

$$\mathbf{prox}_r(z) = \arg\min_x \left\{ r(x) + \frac{1}{2}||x - z||_2^2 \right\}$$

This allows me to use proximal gradient methods to iteratively optimize the objective. Some regularizers I will try are the l_1 norm, the l_2 norm, and a strict bound on the cardinality ($r(x) = 0$ for $\mathbf{card}(x) \leq s$ and ∞ elsewhere), which is useful in situations involving sparse coding. I will compare the accuracy and runtime of these extensions with the ER and HIO algorithms.

4 Implementation

I plan to implement my solution in MatLAB and Python and test it on publicly available high-resolution digital photographs from SCIEN.

References

[CSV13] E. J. Candes, T. Strohmer, and V. Voroninski. PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.

[Fie82] J. R. Fienup. Phase retrieval algorithms: a comparison. *Applied optics*, 21(15):2758–2769, 1982.

[JEH15] K. Jaganathan, Y. C. Eldar, and B. Hassibi. Phase retrieval: an overview of recent developments. *arXiv preprint arXiv:1510.07713*, 2015.

[MAOM13] Y. L. Montagner, E. D. Angelini, and J-C. Olivo-Marin. Phase retrieval with sparsity priors and application to microscopy video reconstruction. In *2013 IEEE 10th International Symposium on Biomedical Imaging*, pages 604–607. IEEE, 2013.

[SWL16] L. Shi, G. Wetzstein, and T. J. Lane. A flexible phase retrieval framework for flux-limited coherent x-ray imaging. *arXiv preprint arXiv:1606.01195*, 2016.