

# HIGH DYNAMIC RANGE IMAGING USING DEEP IMAGE PRIORS

Gauri Jagatap and Chinmay Hegde

Iowa State University

## ABSTRACT

Traditionally, dynamic range enhancement for images has involved a combination of contrast improvement (via gamma correction or histogram equalization) and a denoising operation to reduce the effects of photon noise. More recently, modulo-imaging methods have been introduced for high dynamic range photography to significantly expand dynamic range at the sensing stage itself. The transformation function for both of these problems is highly non-linear, and the image reconstruction procedure is typically non-convex and ill-posed. A popular recent approach is to regularize the above inverse problem via a neural network prior (such as a trained autoencoder), but this requires extensive training over a dataset with thousands of paired regular/HDR image data samples.

In this paper, we introduce a new approach for HDR image reconstruction using neural priors that require no training data. Specifically, we employ deep image priors, which have been successfully used for imaging problems such as denoising, super-resolution, inpainting and compressive sensing with promising performance gains over conventional regularization techniques. In this paper, we consider two different approaches to high dynamic range (HDR) imaging – gamma encoding and modulo encoding – and propose a combination of deep image prior and total variation (TV) regularization for reconstructing low-light images. We demonstrate the significant improvement achieved by both of these approaches as compared to traditional dynamic range enhancement techniques.

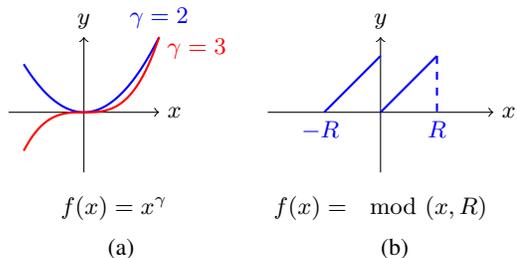
**Index Terms**— Deep image prior, untrained neural networks, convolutional networks, low-light enhancement, HDR imaging, inverse imaging, modulo camera.

## 1. INTRODUCTION

### 1.1. Motivation

Low-light images, or images captured in poor lighting conditions, are ubiquitous in numerous real-world imaging applications. A combination of limited camera sensor hardware and high photon noise can result in low (effective) dynamic range resolution of the captured images. This motivates the problem of low-light image enhancement via extending (post-acquisition) the dynamic range of the captured images, which can be particularly challenging. Several existing low-light enhancement techniques suffer from over-enhancement and the resulting images appear unnatural. Our goal in this paper is to propose novel data processing techniques that can produce improved high dynamic range (HDR) images from acquired camera sensor data.

Mathematically, low-light image acquisition can be viewed as a non-linear forward problem where each “true pixel value” is distorted by different amounts. When a gamma correction model is assumed, the forward transfer function is indicated in Figure 1(a). Low-light images are also corrupted by (additive) photon sensor



**Fig. 1.** Transfer functions for (a) gamma encoding and with factor  $\gamma > 1$  (b) modulo encoding, where  $x$  is the true image and  $y$  are measurements and  $R$  is periodicity of modulo function.

noise, so that the effects of this noise are amplified in a non-linear manner when gamma correction is applied.

An alternative technique for extending dynamic range is to devise new transfer functions with favorable properties. Modulo cameras [1] have been recently proposed to achieve ultra high dynamic range imaging by wrapping the illumination data in a periodic manner. Whenever the pixel value of the camera sensor saturates to its maximum capacity during photon collection, the pixel counter is reset to zero and photon collection continues till the next saturation point (Figure 1(b)). Modulo cameras can theoretically achieve unbounded high dynamic range; however the task of inverting modulo imaging measurements is highly ill-posed.

Deep neural network approaches have led to unprecedented success in solving several ill-posed inverse imaging problems. Image denoising [2], super-resolution [3], inpainting, compressed sensing [4] and phase retrieval [5] are among the several imaging applications that have benefited from the usage of deep convolutional networks (CNNs); such CNN models have slowly replaced hand-crafted priors such as sparsity, total variation and block-matching [4, 6, 7, 8].

However, the main challenge with these approaches is the requirement of massive amounts of training data. In contrast, recently, there has been a surge of interest in using *untrained* neural networks as image priors for regularizing inverse problems. Seminal works on such *deep image priors* [9, 10, 11] are capable of solving linear inverse imaging problems with no training data whatsoever, while merely imposing an auto-encoder [9] or decoder [10, 11] architecture as structural priors on the image. These approaches have been shown to exhibit superior image reconstruction performance as compared to conventional denoising, super-resolution, or compressive sensing approaches [12, 13, 14, 15].

### 1.2. Our contributions

In this paper, extending [11], we introduce a new approach that uses untrained deep neural networks as image priors for dynamic range enhancement of low-light images. We assume that the im-

age  $x^*$  can be (approximately) modeled as a vector belonging to the range spanned by the weights of a deep neural network [9, 10], i.e.,  $x^* = G(\mathbf{w}; z)$  where  $\mathbf{w}$  is a set of the weights of the deep network,  $z$  is a fixed latent code, and  $G$  captures the overall deep network architecture. We incorporate this prior into natural iterative regularization algorithms: (a) one that simultaneously corrects for brightness (which is modeled as a gamma correction) and denoises a low-light image, and (b) one that reconstructs HDR images from modulo measurements<sup>1</sup>. Our experimental results demonstrate significant improvement over conventional regularized inverse-imaging methods with hand-crafted, while avoiding the need for any training data whatsoever.

### 1.3. Prior work

Due to space limitations we only provide a brief overview of prior work. An extensive review of low-light image enhancement schemes can be found in [16].

Low light image enhancement has been extensively studied in the image processing literature. Statistical image reconstruction methods include Histogram Equalization (HE) which involves adjusting the pixel intensities such that the histogram is flat and uniformly spread across all possible values instead of concentrated near low pixel intensities. Further, different regularizations may be imposed on the histogram statistics (for example, contextual and variational contrast enhancement (CVC) [17] and local histogram equalization (HE) [18, 19] which adopt the sliding window strategy to perform HE locally). Another approach is to use an appropriate transformation model such as (adaptive) gamma correction, which can correlate pairs of bright and dark images [20]. Concurrently, retinex theory [21] based approaches assumes that the amount of light reaching observers can be decomposed into two parts: illumination and scene reflection [22].

Instead of formulating an appropriate forward model based on the imaging physics, a mapping may also be *learned* by training a neural network over thousands of example images [23]. LL-Net [23] uses an autoencoder based approach which emulates a gamma correction and denoising model trained on pairs of artificially darkened and corrupted images, whereas [24] introduces a new neural network model called Deep-Retinex. EnlightenGAN [25] uses a CycleGAN [26] based approach for image to image translation. The merits of such approaches is that after it is trained, the low-light correction operation is a one-shot procedure. The trained model implicitly learns an appropriate scaling of pixel values, and applies an appropriate regularization based on image database statistics. There are several shortcomings of neural network-based approaches. First, all of these approaches require a large database of paired [24, 23] or unpaired [25] bright and dark images. The requirement to train a huge neural network with thousands of parameters can incur unreasonably high computational costs. Second, this approach allows little control on the *degree* of correction or enhancement for a given image.

The above alleviate the problem of photon saturation, modulo-cameras were introduced in [1]. This computational sensing technique allows for infinite dynamic range [27]. However the inverse mapping is highly ill-posed and requires computational expensive approaches for signal reconstruction based on graph cuts [1] or quadratically constrained quadratic program [28]. [29] uses a sparsity prior to develop an alternating minimization technique MoRAM, to solve this problem from compressive Gaussian measurements.

<sup>1</sup>For the low-light image enhancement problem, we additionally assume that the image exhibits low total variation (TV).

### 1.4. Notation

Throughout the paper, lower case letters denote vectors, such as  $v$  and upper case letters for matrices, such as  $M$ . The indicator vector is denoted by  $\mathbf{1}_{\text{cond}(v)}$  and is a boolean vector the size of  $v$  and indicates element-wise satisfaction of a given condition  $\text{cond}(\cdot)$ . The set of weights of a neural network is represented by  $\mathbf{w} = \{W_1, \dots, W_L\}$  where  $L$  denotes the number of layers.

## 2. PROBLEM SETUP

In this section we describe two image formation models and two corresponding HDR image reconstruction techniques that employ deep image priors.

### 2.1. Low-light model

Our first method reconstructs HDR images (assumed an underlying deep image prior) from noisy dark or low-light (LL) images  $y \in \mathbb{R}^d$  that are modeled by a gamma correction model that obeys Steven’s law corrupted by additive white Gaussian noise  $\varepsilon$  with variance  $\sigma$ :

$$f(x^*) = y = c \cdot x^{*\gamma} + \varepsilon, \quad (1)$$

where  $x^* \in \mathbb{R}^d$  is the true image (whose pixel values lie between 0 and 1), factor  $\gamma \in \mathbb{R}^+$  is such that  $\gamma > 1$  corresponds to a darkened observed image  $y$  and  $0 < \gamma < 1$  corresponds to a brightened observed image, and  $c \in \mathbb{R}^+$  is a known scaling constant. Since we consider the problem of low-light enhancement, for our problem setting  $\gamma > 1$  is known.

### 2.2. Modulo observation model

Our second method reconstructs HDR images from modulo-valued observations  $y \in \mathbb{R}^n$ , with periodicity  $R \in \mathbb{R}^+$ :

$$f(x^*) = y = \text{mod}(Ax^*, R)$$

where  $x^* \in \mathbb{R}^d$  is the true image, with pixel values lying between 0 and 1. Here  $A \in \mathbb{R}^{n \times d}$  is a linear transformation matrix which models optical acquisition setup. In this paper, we consider a model with two periods, following [29]. Mathematically:

$$f(x) = y = Ax^* + R \cdot \mathbf{1}_{Ax^* < R} \quad (2)$$

where  $\mathbf{1}_{Ax^* < R}$  is an element-wise indicator.

### 2.3. Deep image prior

A deep image prior assumes that the image  $x \in \mathbb{R}^{d \times k}$  ( $k$  channels each of dimension  $d$ ) can be represented as the action of a deep generative network  $G(\mathbf{w}; z)$  with unknown weights  $\mathbf{w}$  on some fixed latent code  $z$ , such that  $x = G(\mathbf{w}; z)$ . The latent code  $z := \text{vec}(Z_1)$  with  $Z_1 \in \mathbb{R}^{d_1 \times k_1}$  is a low-dimensional parameter with dimension  $d_1 k_1 \ll dk$ ; its elements are generated from an appropriately scaled uniform random distribution and are kept fixed throughout.

The task is to estimate the image

$$\hat{x} = G(\hat{\mathbf{w}}; z) \approx G(\mathbf{w}^*; z) = x^*$$

and corresponding weights  $\hat{\mathbf{w}}$ , for a fixed seed  $z$ , where  $x^*$  is assumed to be the true image and the true weights  $\mathbf{w}^*$  (possibly non-unique) satisfy

$$\mathbf{w}^* = \min_{\mathbf{w}} \|x^* - G(\mathbf{w}; z)\|_2^2.$$

---

**Algorithm 1** Low-light image enhancement with Deep Image Prior.

---

- 1: **Input:**  $A, z = \text{vec}(Z_1), \eta, \mathbf{w}^0$ .
  - 2: **while** termination condition not met **do**
  - 3:    $\mathbf{w}^{t+1} \leftarrow \mathbf{w}^t - \eta \nabla \mathcal{L}(\mathbf{w}^t)$                     {gradient step}
  - 4: **end while**
  - 5: **Output**  $\hat{x} \leftarrow G(\mathbf{w}^T; z)$ .
- 

Substituting the surjective mapping  $G : \mathbf{w} \rightarrow x$ , and optimizing over  $\mathbf{w}$ , we have

$$\min_{\mathbf{w}} \mathcal{L}(\mathbf{w}) := \min_{\mathbf{w}} \|y - f(G(\mathbf{w}; z))\|_2^2 + \lambda \cdot \text{TV}(G(\mathbf{w}; z)), \quad (3)$$

to obtain  $\hat{\mathbf{w}} = \arg \min_{\mathbf{w}} \mathcal{L}(\mathbf{w})$  and corresponding image  $\hat{x}$ . Here, we have imposed a Mean-Squared Error (MSE) loss, and an additional Total Variation (TV) regularization to smoothen out the effective image  $\hat{x} = G(\hat{\mathbf{w}}; z)$  with Lagrangian multiplier  $\lambda$ , which is chosen appropriately. Alternatively, one can use a variable splitting approach of the form,

$$\begin{aligned} \min_{x, \mathbf{w}} \mathcal{L}(x, \mathbf{w}) &:= \min_x \|y - f(x)\|_2^2 + \lambda \cdot \text{TV}(x), & (4) \\ \text{s.t. } &x = G(\mathbf{w}; z) \end{aligned}$$

Specifically, the untrained network  $G(\mathbf{w}; z)$  takes the form of an expansive neural network; a decoder architecture similar to the one in [10], or that of the generator of a DCGAN [30, 31]. Fixing this architectural framework, we define the deep network prior as follows:

**Definition 1.** A given image  $x \in \mathbb{R}^{d \times k}$  is said to obey an untrained neural network prior if it belongs to a set  $\mathcal{S}$  defined as:  $\mathcal{S} := \{x | x = G(\mathbf{w}; z)\}$ , where  $z$  is a (randomly chosen, fixed) latent code vector and  $G(\mathbf{w}; z)$  has the form of a generator [30] or decoder [10].

### 3. RECONSTRUCTION ALGORITHMS

#### 3.1. Low-light model

A combination of the forward model described above coupled with deep image prior leads to a natural recovery algorithm. The loss function in Eq. 3 with  $f$  assuming the form in Eq. 1 can be (heuristically) minimized using gradient descent by optimizing over the weights  $\mathbf{w}$  of the deep image prior; the procedure is described in Algorithm 1. The gradients  $\nabla \mathcal{L}(\mathbf{w})$  in Step 3 of Algorithm 1 are computed using back-propagation using standard software frameworks for deep learning such as TensorFlow and PyTorch, with learning rate  $\eta$ . The weights  $\mathbf{w}^0$  are initialized according to He’s initialization [32]. Each iteration  $t$  yields a new image estimate  $x^t = G(\mathbf{w}^t; z)$  which belongs to the range of the deep image prior,  $\mathcal{S}$ .

#### 3.2. Modulo observation model

The forward model in Eq. 2 is both non-differentiable and non-bijective, so we require additional algorithmic adjustments. In Algorithm 2, we propose a two-stage reconstruction approach.

The basic idea is to first estimate the “bin” into which the modulo observations are folded; if we can do this well, then it essentially becomes a linear inverse problem. Let  $p$  denote the bin-index of vector  $Ax$  which is ‘1’ if  $Ax < 0$  and ‘0’ if  $Ax > 0$ . Then, a crude estimate of  $p = p_{init}$  can be obtained as follows, as per [29]:

$$p_{init} = \begin{cases} 0, & y < \frac{R}{2} \\ 1, & y > \frac{R}{2} \end{cases}. \quad (5)$$

---

**Algorithm 2** Modulo sensing with Deep Image Prior.

---

#### INITIALIZATION STAGE

---

- 1: **Input:**  $A, z = \text{vec}(Z_1), \eta, \mathbf{w}^0$ .
  - 2:  $y_{init} = y - R \cdot p_{init}$
  - 3:  $x^0 = \arg \min_{x \in \mathcal{S}} \|y_{init} - Ax\|_2^2$  {Net-GD for CS}
- 

#### DESCENT STAGE

---

- 4: **Input:**  $A, z, x^0, \eta, \mathbf{w}^0$
  - 5: **while** termination condition not met **do**
  - 6:    $p^t = \mathbf{1}_{Ax^t < 0}$
  - 7:    $y_c = y - R \cdot p^t$
  - 8:    $v^t \leftarrow x^t - \eta \nabla_x \|y_c - Ax\|_2^2$                     {gradient step}
  - 9:    $x^{t+1} \leftarrow \arg \min_{x \in \mathcal{S}} \|v^t - x\|_2^2$  {project to  $\mathcal{S}$ }
  - 10: **end while**
  - 11: **Output**  $\hat{x} \leftarrow G(\mathbf{w}^T; z)$ .
- 

With  $p_{init}$ , one can first solve a linear recovery problem from imperfect measurements  $y_{init} = y - R \cdot p_{init}$ , (Steps 2 and 3 of Algorithm 2), subject to the constraint that the reconstructed signal  $x^0$  lies in the range of a deep image prior. This can be implemented using the Net-GD algorithm from [11]. In our experiments, we observed that this procedure yields a good enough initialization for the problem of signal reconstruction from modulo measurements.

Given a good initialization, one can apply an alternating minimization type scheme as described in Steps 6-9 of Algorithm 2; these alternately refine the estimates of the image as well as the bin-indices, mirroring the approach of [33, 34]. We formulate the ‘corrected’ linear measurements  $y_c = y - R \cdot p^t$  and use these to solve a linear recovery problem with a single step of projected gradient descent in Steps 8 and 9. This is similar to solving the setting in Eq. 4, but with a warm start (and with  $\lambda = 0$ ). For both stages of Algorithm 2, we initialize the weights  $\mathbf{w}^0$  as per He initialization [32]. The gradients  $\nabla_x \|y_c - Ax\|_2^2$  in Step 8 of Algorithm 2 can be computed using back-propagation with learning rate  $\eta$ . The projection in Step 9 of Algorithm 2 can be solved as  $\min_{\mathbf{w}} \|v^t - G(\mathbf{w}; z)\|_2^2$  via the Adam optimizer.

### 4. EXPERIMENTAL RESULTS

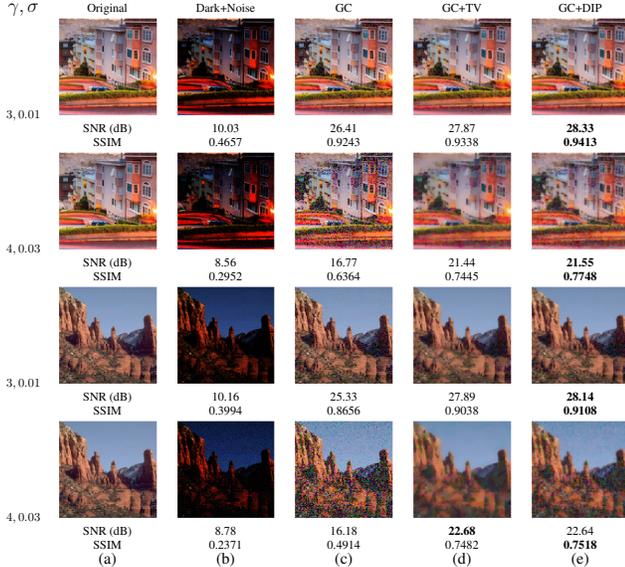
**Performance metrics:** We compare reconstruction quality using two metrics: (i) normalized Mean-Squared Error (nMSE), which is calculated as  $\|\hat{x} - x^*\|^2 / \|x^*\|^2$ , and (ii) Structural Similarity Index (SSIM) [35]. We tabulate their variation against different modeling parameters such as noise level  $\sigma$ ,  $\gamma$  (for low-light image enhancement) and different compression rates  $f = n/d$  (for modulo signal recovery), averaged over all trials.

#### 4.1. Low-light model

**Forward model:** We select  $\gamma = 3, 4$  and noise variance  $\sigma = 0.01, 0.03$ . For sake of simplicity we fix scaling constant  $c = 1$ .

**Dataset:** We use two test images: San Francisco (rows 1 and 2 of Figure 2) and Red Rock (rows 3 and 4 of Figure 2). Both have been reshaped to  $128 \times 128 \times 3$  ( $d = 49, 152$ ). The pixel values of all images are scaled to lie between 0 and 1.

**Deep network architecture:** We use a 4 layer neural network architecture with a DCGAN architecture consisting of transposed convolutions, ReLU operations and batch normalizations. The final



**Fig. 2.** (LL) (a) Original image, (b) image darkened with factor  $\gamma$ , followed by addition of noise with variance  $\sigma$ , (c) gamma corrected image (d) gamma correction followed by TV denoising (e) Algorithm 1 using Deep Image Prior.

layer has a sigmoid activation, to smoothen out the output image. The widths of each layer are varied between  $k = 40$  to  $k = 100$  and the performance corresponding to the best setting is reported in Figure 2. Further details on implementation as well as additional experiments can be found in the supplementary material [36].

**Algorithms and baselines:** We implement 3 schemes for the low-light image enhancement problem - (i) gamma correction (ii) gamma correction followed by denoising via Total Variation (TV) denoising and (iii) deep image prior based denoising (Algorithm 1).

**Implementation details:** Algorithm 1 was implemented using the PyTorch framework with Python 3, GPU support, using Adam optimizer. For implementing TV denoising, we use the sci-kit image library. Lagrangian parameter  $\lambda$  is tuned for each image appropriately.

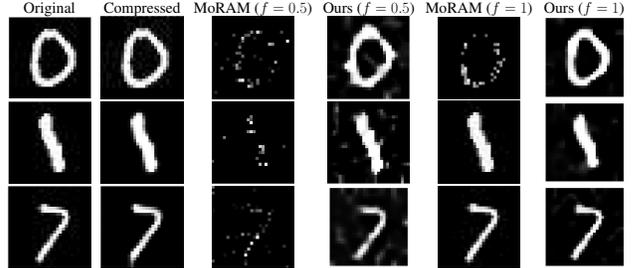
**Running time:** We report the average running times for different algorithms across different  $\gamma$  and  $\sigma$  values for our example images is 35.43s (Algorithm 1), and 0.16s (TV).

From our results in Figure 2, we note that DIP prior coupled with TV regularization can perform comparable or better image enhancement from darkened and corrupted images.

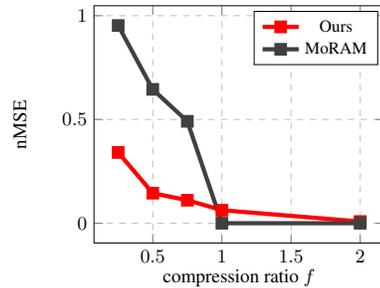
## 4.2. Modulo measurements

**Dataset:** We use 3 images from the MNIST database to test our algorithms and reconstruct 3 grayscale images each of size  $28 \times 28$  pixels ( $d = 784$ ). The pixel values of all images are scaled to lie between 0 and 1.

**Forward model:** We use a Gaussian measurement matrix of size  $n \times d$  with  $n$  varied such that  $n/d = 0.25, 0.5, 0.75, 1, 2$ . The elements of  $A$  are picked such that  $A_{i,j} \sim \mathcal{N}(0, 1/n)$  and we report averaged reconstruction error values over 10 different instantiations of  $A$  for a fixed image (image of digit ‘1’ from MNIST), network configuration and compression ratio  $n/d$  in Figure 3 (b). The period  $R$  is picked as  $R = (1.8) \cdot \max_i a_i^T x^*$ , where  $i = \{1, \dots, n\}$ .



(a)



(b)

**Fig. 3.** (Modulo) Reconstructed images from modulo measurements (a) at compression rates of  $f = n/d = 0.5, 1$  for MNIST images, (b) nMSE at different compression rates  $f = n/d$  for MNIST digit ‘1’ averaged over 10 trials.

**Deep network architecture:** We use the deep decoder architecture [10] with a 3 layer configuration with channel sizes  $k_1 = 25, k_2 = 15, k_3 = 10$  and consisting of ReLU, upsampling and  $1 \times 1$  convolutions  $W_1 \in \mathbb{R}^{k_1 \times k_2}, W_2 \in \mathbb{R}^{k_2 \times k_3}$  and  $W_3 \in \mathbb{R}^{k_3 \times 1}$ . We further use batch normalization on all layers and a sigmoid operation on the final layer. For further details refer to the setup in [11].

**Algorithms and baselines:** We implement 2 schemes for signal reconstruction from modulo measurements, (i) Algorithm 1 and (ii) Modulo recovery using Alternating Minimization (MoRAM) [29] for all images.

**Implementation details:** Algorithm 2 was implemented using PyTorch framework with Python 3, using GPU support. Adam optimizer was used for solving Steps 3 and 9 of Algorithm 2 whereas SGD optimizer was used to compute single step of update in Step 8. MoRAM [29] was implemented using the open-source MATLAB code, while using CoSamP [37] and Justice Pursuit [38] interchangeably for the sparse signal estimation stage, depending on which approach yielded better performance for the images chosen. Since the images chosen from MNIST are sparse to begin with, we did not use any sparsifying transform on the images.

**Running time:** We report the average running times for different algorithms across different measurement levels for examples from MNIST is 25.59s (Algorithm 2), 9.1s (MoRAM).

From our results in Figure 3, we can see that DIP prior can reconstruct images from modulo images at much lower compression rates. At higher sample rates, sparsity priors are more beneficial due to the fact that we have chosen an under-parameterized network for our prior which corresponds to small representational error.

## 5. REFERENCES

- [1] H. Zhao, B. Shi, C. Fernandez-Cull, S. Yeung, and R. Raskar, "Unbounded high dynamic range photography using a modulo camera," in *2015 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 2015, pp. 1–10.
- [2] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P. Manzagol, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *Journal of machine learning research*, vol. 11, no. Dec, pp. 3371–3408, 2010.
- [3] C. Dong, C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 2, pp. 295–307, 2016.
- [4] J. Chang, C. Li, B. Póczos, and B. Kumar, "One network to solve them all—solving linear inverse problems using deep projection models," in *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE, 2017, pp. 5889–5898.
- [5] C. Metzler, P. Schniter, A. Veeraraghavan, and R. Baraniuk, "prdeep: Robust phase retrieval with a flexible deep network," in *International Conference on Machine Learning*, 2018, pp. 3498–3507.
- [6] A. Bora, A. Jalal, E. Price, and A. Dimakis, "Compressed sensing using generative models," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 537–546.
- [7] P. Hand, O. Leong, and V. Voroninski, "Phase retrieval under a generative prior," in *Advances in Neural Information Processing Systems*, 2018, pp. 9136–9146.
- [8] T. Lillicrap Y. Wu, M. Rosca, "Deep compressed sensing," *arXiv preprint arXiv:1905.06723*, 2019.
- [9] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9446–9454.
- [10] R. Heckel and P. Hand, "Deep decoder: Concise image representations from untrained non-convolutional networks," in *International Conference on Learning Representations*, 2018.
- [11] G. Jagatap and C. Hegde, "Algorithmic guarantees for inverse imaging with untrained network priors," *arXiv preprint arXiv:1906.08763*, 2019.
- [12] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM review*, vol. 43, no. 1, pp. 129–159, 2001.
- [13] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising with block-matching and 3d filtering," in *Image Processing: Algorithms and Systems, Neural Networks, and Machine Learning*. International Society for Optics and Photonics, 2006, vol. 6064, p. 606414.
- [14] V. Pappas, Y. Romano, J. Sulam, and M. Elad, "Convolutional dictionary learning via local processing," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5296–5304.
- [15] C. Li, W. Yin, and Y. Zhang, "User's guide for tval3: Tv minimization by augmented lagrangian and alternating direction algorithms," .
- [16] Yuen Peng Loh, Xuefeng Liang, and Chee Seng Chan, "Low-light image enhancement using gaussian process for features retrieval," *Signal Processing: Image Communication*, vol. 74, pp. 175–190, 2019.
- [17] T. Celik and T. Tjahjadi, "Contextual and variational contrast enhancement," *IEEE Transactions on Image Processing*, vol. 20, no. 12, pp. 3431–3441, 2011.
- [18] Ali M Reza, "Realization of the contrast limited adaptive histogram equalization (clahe) for real-time image enhancement," *Journal of VLSI signal processing systems for signal, image and video technology*, vol. 38, no. 1, pp. 35–44, 2004.
- [19] S. Pizer, E. P. Amburn, J. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. ter Haar Romeny, J. Zimmerman, and K. Zuiderveld, "Adaptive histogram equalization and its variations," *Computer vision, graphics, and image processing*, vol. 39, no. 3, pp. 355–368, 1987.
- [20] S. Huang, F. Cheng, and Y. Chiu, "Efficient contrast enhancement using adaptive gamma correction with weighting distribution," *IEEE transactions on image processing*, vol. 22, no. 3, pp. 1032–1041, 2012.
- [21] E. Land, "The retinex theory of color vision," *Scientific american*, vol. 237, no. 6, pp. 108–129, 1977.
- [22] X. Guo, Y. Li, and H. Ling, "Lime: Low-light image enhancement via illumination map estimation," *IEEE Transactions on image processing*, vol. 26, no. 2, pp. 982–993, 2016.
- [23] K. Lore, A. Akintayo, and S. Sarkar, "Llnet: A deep autoencoder approach to natural low-light image enhancement," *Pattern Recognition*, vol. 61, pp. 650–662, 2017.
- [24] C. Wei, W. Wang, W. Yang, and J. Liu, "Deep retinex decomposition for low-light enhancement," *arXiv preprint arXiv:1808.04560*, 2018.
- [25] Y. Jiang, X. Gong, D. Liu, Y. Cheng, C. Fang, X. Shen, J. Yang, P. Zhou, and Z. Wang, "Enlightengan: Deep light enhancement without paired supervision," *arXiv preprint arXiv:1906.06972*, 2019.
- [26] J. Zhu, T. Park, P. Isola, and A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [27] A. Bhandari, F. Krahmer, and R. Raskar, "Unlimited sampling of sparse signals," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4569–4573.
- [28] M. Cucuringu and H. Tyagi, "On denoising modulo 1 samples of a function," in *International Conference on Artificial Intelligence and Statistics*, 2018, pp. 1868–1876.
- [29] V. Shah and C. Hegde, "Signal reconstruction from modulo observations," *arXiv preprint arXiv:1812.00557*, 2018.
- [30] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [31] D. Van Veen, A. Jalal, E. Price, S. Vishwanath, and A. Dimakis, "Compressed sensing with deep image prior and learned regularization," *arXiv preprint arXiv:1806.06438*, 2018.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
- [33] G. Jagatap and C. Hegde, "Fast, sample-efficient algorithms for structured phase retrieval," in *Advances in Neural Information Processing Systems*, 2017, pp. 4917–4927.
- [34] G. Jagatap and C. Hegde, "Sample-efficient algorithms for recovering structured signals from magnitude-only measurements," *IEEE Transactions on Information Theory*, 2019.
- [35] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [36] G. Jagatap and C. Hegde, "High dynamic range imaging using deep image priors," <https://gaurijagatap.github.io/assets/hdrimage.pdf/>, 2019.
- [37] D. Needell and J. Tropp, "Cosamp: Iterative signal recovery from incomplete and inaccurate samples," *Applied and computational harmonic analysis*, vol. 26, no. 3, pp. 301–321, 2009.
- [38] J. Laska, M. Davenport, and R. Baraniuk, "Exact signal recovery from sparsely corrupted measurements through the pursuit of justice," in *2009 Conference Record of the Forty-Third Asilomar Conference on Signals, Systems and Computers*. IEEE, 2009, pp. 1556–1560.