Non-Local Compressive Sampling Recovery*

Xianbiao Shu[1], Jianchao Yang[2] and Narendra Ahuja[1]

[1]University of Illinois at Urbana-Champaign, [2] Adobe Research Lab

[1]{xshu2, n-ahuja}@illinois.edu, [2]jiayang@adobe.com

Abstract

Compressive sampling (CS) aims at acquiring a signal at a sampling rate below the Nyquist rate by exploiting prior knowledge that a signal is sparse or correlated in some domain. Despite the remarkable progress in the theory of CS, the sampling rate on a single image required by CS is still very high in practice. In this paper, a non-local compressive sampling (NLCS) recovery method is proposed to further reduce the sampling rate by exploiting non-local patch correlation and local piecewise smoothness present in natural images. Two non-local sparsity measures, i.e., non-local wavelet sparsity and non-local joint sparsity, are proposed to exploit the patch correlation in NLCS. An efficient iterative algorithm is developed to solve the NLCS recovery problem, which is shown to have stable convergence behavior in experiments. The experimental results show that our NLCS significantly improves the state-of-the-art of image compressive sampling.

1. Introduction

We have been witnessing the rapid development of digital image sensors with ever-increasing fidelity and resolution. Conventional digital sensors follow Shannon's Nyquist sampling theorem, which requires that the sampling rate be above the Nyquist rate, i.e., twice the maximal analog signal frequency. While incurring no loss of information, Nyquist sampling generates a large amount of raw data, which is challenging to acquire, encode and transfer in many applications such as infrared imaging, magnetic resonance imaging (MRI) and wireless sensor networks.

Recently, compressive sensing [8] or compressive sampling (CS) [5], has been developed to reduce the sampling rate below the Nyquist rate. Its main idea is that a signal can be decoded from incomplete compressive measurements by seeking its sparsity in some domain. The resulting sampling rate (defined as the ratio of the sample count to the signal size) is roughly proportional to the signal sparsity. Much effort has been made to further reduce the sampling rate of CS by exploring prior knowledge of natural images and videos. The state-of-the-art method in image CS (2DCS) [9, 14, 15] exploits two kinds of prior knowledge of natural images/videos—piecewise smoothness by total variation (TV) [18] and sparsity in the 2D wavelet domain. With this prior knowledge, it recovers an image I from its random measurements B as follows:

$$\min_I \mathrm{TV}(I) + \lambda \|\Psi_{2D}(I)\|_1 \quad \text{s. t.} \quad \Phi I = B, \qquad (1)$$

where Φ is the sampling matrix, Ψ_{2D} denotes the 2D wavelet transform and λ is a regularization constant. However, due to that natural images are not sufficiently sparse in the wavelet domain, 2DCS still requires a high sampling rate and its recovery tends to blur sharp edges and texture. Thus, the full promise of image CS remains unrealized.

In this paper, we propose a non-local compressive sampling (NLCS) recovery method, which further reduces the sampling rate of image CS by exploiting non-local

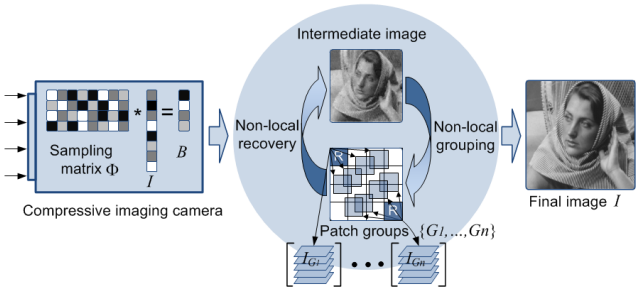

Figure 1. Imaging via non-local compressive sampling (NLCS), which consists of two steps—sensing and recovery. First, it acquires the compressive measurement B by a random sampling matrix Φ. Second, NLCS recovers an image I from the measurements $B = \Phi I$ by iterating between two steps—non-local grouping and non-local recovery.

*The support of the Office of Naval Research under grant N00014-12-1-0259 is gratefully acknowledged.

patch correlation and the conventional piecewise smoothness prior in natural images. It is mainly motivated by the recent advances in non-local mean approaches [4, 16, 7] in image restoration. Non-local mean approaches successfully achieve the state-of-the-art performance in image restoration [7] by seeking the correlation of image patches. As a common prior in natural images, the patch correlation should help reduce the required sampling rate in CS.

Different from [11] that enforeces patch correlation by recursive filtering, we enforce the patch correlation by a non-local sparsity (NLS) measure and combine this measure with the conventional TV measure in a neat objective function. Two kinds of NLS measures, i.e., non-local wavelet sparsity (NLWS) and non-local joint sparsity (NLJS), are proposed, both of which can capture the patch correlation, as well as the wavelet sparsity of patches themselves. Compared with NLWS used in image denoising [7], NLJS better enforeces the patch correlation and thus achieves higher recovery accuracy in CS.

Motivated by BM3D [7] in image denoising, we explore the non-local sparsity of an image by clustering the image patches into multiple groups and then imposing the correlation prior within each group. However, in contrast with the image denoising application, the original image is unknown in CS except for its random measurements B, and thus the patch-grouping information is unavailable prior to the NLCS recovery. Thus, our NLCS needs to recover both the unknown image and the grouping information about its patches, which turns out to be a chicken-and-egg problem[1].

In this paper, we address this problem by an iterative scheme, where the unknown image and its patch-grouping information are estimated alternatively. Figure 1 gives the flow chart of our NLCS. First, the compressive measurement B is obtained from a scene by a random sampling matrix Φ . Then, our NLCS recovers an image I from the measurement $B = \Phi I$ in two iterative steps: (1) given an intermediate image I, non-local grouping clusters its patches into groups $\{G_i\}_{i=1}^n$; (2) Given $\{G_i\}_{i=1}^n$, non-local recovery estimates an image I from B. An efficient algorithm is proposed to solve this non-local recovery problem.

Related Work. Some research work has been done on using patch or frame correlation for compressive sensing. But, the way they use it is significantly different from our NLCS. [11] proposed a non-parametric CS method, which recursively uses non-local denoising filter (BM3D) to fill the unobserved portion of the Fourier spectrum. This is quite different from our parametric method—NLCS, which has an explicit objective function including non-local sparsity measure and piecewise smoothness. In addition, this method only works for Radon or Fourier sampling while our NLCS is suitable for general CS applications. [10] proposed a learning based CS method, which jointly designs and optimizes the patch-sensing matrix and the over-complete patch dictionary. In contrast with this patch-wise method, our approach is a holistic CS method that takes measurements on the entire image and recovers the image by taking advantage of its self-similarity prior. It is worthwhile mentioning most existing CS systems, e.g., MRI modalities and single-pixel camera [9], use holistic measurements. Also, this dictionary learning method cannot provide a generic solution to CS, since its performance largely depends on the training dataset. Three-dimensional CS (3DCS) [20] significantly reduces the sampling rate of video CS by exploiting the temporal correlation (low-rank) of a video. This motivates us to explore the patch correlation in image CS. Compared with seeking the temporal correlation, where the video frames are highly corrected, the patch correlation is more challenging to exploit in image CS, due to the fact that the patch-grouping information is unknown in CS.

[1]Estimating one component requires and also benefits from the other.

2. Non-Local Compressive Sampling (NLCS)

2.1. Formulation

In additional to the piecewise smoothness prior and wavelet sparsity used in 2DCS, our non-local compressive sampling (NLCS) recovery exploits another ubiquitous prior knowledge about natural images—patch correlation. Let $I \in \mathbb{R}^{M\times N}$ denote a 2D image, and I_x a $d \times d$ image patch at location x (2D coordinate) on I. The patches of image I are divided into n groups $\mathcal{G} = \{G_1, ..., G_n\}$ by non-local grouping based on some similarity measure, where G_i contains the coordinates of the patches belonging to the i-th group. Patches from the i-th group stack into a 3D cube denoted by $I_{G_i} \in \mathbb{R}^{d\times d\times m_i}$, where m_i is the number of patches in group G_i. Based on the grouping information $\mathcal{G}$, we seek patch correlation of image I by minimizing its non-local sparsity. This non-local sparsity imposes the requirements that (1) image patches repeat themselves across the image and (2) image patches are sparse in some domain. Taking advantage of this non-local sparsity of images, our NLCS recovers a sharp and piecewise smooth image by

$$\min_{I,\mathcal{G}} \text{TV}_{\ell_1}(I) + \lambda \text{NLS}(I, \mathcal{G}) \quad \text{s. t.} \quad \Phi I = B, \qquad (2)$$

where $\text{NLS}(I, \mathcal{G})$ measures the non-local sparsity of the image based on the patch grouping information $\mathcal{G}$. Here, we employ the ℓ_1-norm based total variation $\text{TV}_{\ell_1}(I) = \|D_1 I\|_1 + \|D_2 I\|_1$, due to its superiority to the traditional total variation $\text{TV}_{\ell_1\ell_2}$ [19], where D_1 and D_2 are finite difference operators along horizontal and vertical axes. Althought our NLCS is compatible with any compressive sampling matrix, for computational efficieny, we choose the circulant sampling in [17, 20], i.e., $\Phi = SC$, where C is a circulant matrix and S is a random subsampling matrix.

At the core of NLCS is the non-local sparsity measure that gauges the correlation between patches and the sparsity of patches themselves. In this section, we will present two non-local sparsity measures, namely, non-local wavelet sparsity and non-local joint sparsity.

2.2. Non-Local Wavelet Sparsity

Many methods have been proposed to obtain the correlation of a signal ensemble, e.g., imposing its low-rank [6] and learning its low-dimensional subspace [2]. In NLCS, we can adapt these methods to impose the correlation of each patch group I_{G_i} (Figure 1) that is highly correlated and almost lies in a rank-1 subspace. However, these methods fail to take into account another important prior knowledge—the sparsity of patches themselves.

A good candidate of the non-local sparsity measure is the 3D wavelet sparsity. It is popularly used to seek patch correlation in image denoising (e.g., BM3D) and also employed to impose temporal correlation in video CS [21]. Here, we employ it to explore the non-local patch correlation in image CS. On the 2D wavelet coefficients of each patch in a group $\Psi_{2D}(I_{G_i})$, we conduct wavelet transform along the third axis, and compute non-local wavelet sparsity (NLWS) as the sum of the ℓ_1 norm of all 3D wavelet coefficients:

$$\text{NLWS}(I, \mathcal{G}) = \sum_{i=1}^{n} \|\Psi_{3D}(I_{G_i})\|_1, \tag{3}$$

where Ψ_{3D} is the 3D wavelet transform.

2.3. Non-Local Joint Sparsity

Motivated by the observation that the matched patches in one group are almost identical (up to some sparse errors), we propose another non-local sparsity measure—non-local joint sparsity (NLJS). Joint sparsity is initially proposed in [1] to identify the common component and sparse innovation components in a signal ensemble. Here, we apply the joint sparsity idea to a single image I consisting of non-local patch groups ($\mathcal{G}$) and define the non-local joint sparsity of image I as follows:

$$\text{NLJS}(I, \mathcal{G}) = \sum_{i=1}^{n} \text{JS}(I_{G_i}), \tag{4}$$

where the joint sparsity $\text{JS}(I_{G_i})$ is defined as follows. For each patch group I_{G_i}, we first conduct 2D wavelet transform on each patch to obtain $\Psi_{2D}(I_{G_i}) \in \mathbb{R}^{d\times d\times m_i}$. Then, we decompose $\Psi_{2D}(I_{G_i})$ into the sum of a replica of the common component $\overline{Z}_i \in \mathbb{R}^{d\times d}$ and the sparse innovation components $\widehat{Z}_i \in \mathbb{R}^{d\times d\times m_i}$. Then, the joint sparsity of I_{G_i} is defined as

$$\begin{aligned} \text{JS}(I_{G_i}) = \min_{\overline{Z}_i, Z_i} \|\overline{Z}_i\|_1 + \eta\|Z_i\|_1, \\ \text{s.t.} \quad [\overline{Z}_i : ... : \overline{Z}_i] + Z_i = \Psi_{2D}(I_{G_i}), \end{aligned} \tag{5}$$

where $[\overline{Z}_i : ... : \overline{Z}_i]$ denotes a 3D cube consisting of m_i replicas of $\overline{Z}_i$ and η is a regularization constant that is equal to or larger than 1. Minimizing the first term $\|\overline{Z}_i\|_1$ imposes the wavelet sparsity of the patches themselves, similar to 2DCS, while minimizing the second term $\|\widehat{Z}_i\|_1$ imposes the patch correlation within each group. Therefore, for images that contain abundantly repeating local structures, η can be set larger to emphasize the patch correlation prior. Note that we have an explicit solution to the aforementioned minimization problem. The optimal common component $\overline{Z}_i$ can be obtained by elementwise applying the median filter on $\Psi_{2D}(I_{G_i})$ along the third dimension[2].

In sum, NLWS and NLJS are both defined in the 2D wavelet domain, but differ in the ways they impose patch correlation. In NLWS, we conduct 1D wavelet transform on $\Psi_{2D}(I_{G_i})$ and sum up the ℓ_1 norm of all 3D wavelet coefficients. Minimizing NLWS tends to impose the patch correlation and also blur the patches, since all the wavelet coefficients are uniformly penalized. In NLJS, by assuming $\Psi_{2D}(I_{G_i})$ is approximated by the rank-1 subspace (spanned by $\overline{Z}_i$) , we sum up the ℓ_1 norm of $\overline{Z}_i$ and the approximation error $\widehat{Z}_i$. In this way, NLJS can heavily penalize the high-frequency components $\widehat{Z}_i$ while slightly regularizing the low-frequency component $\overline{Z}_i$. Thus, NLJS better retains the sharp edges and textures in the recovered image.

3. An Efficient Algorithm for NLCS

Our NLCS algorithm in Eq. (2) attempts to recover the underlying image I from its compressive measurements ΦI by minimizing its total variation and non-local sparsity. However, the non-local sparsity is defined based on the patch grouping $\mathcal{G}$, which in turn requires knowledge of the image I. Direct minimization over unknown I and $\mathcal{G}$ is intractable. Instead, we present an iterative algorithm to find an approximate solution. The iterative algorithm starts with an initial estimate I recovered by 2DCS, and then iterates between two steps—(1) non-local grouping that extracts the patch grouping $\mathcal{G}$ from image I and (2) non-local joint recovery that recovers the image I based on the updated grouping information $\mathcal{G}$.

3.1. Non-Local Patch Grouping

We use the same block-matching scheme in BM3D [7] for non-local patch grouping. Given an estimated image I, it first obtains n reference patches, denoted as $I_{x_i^r} \in \mathbb{R}^{d\times d}, 1 \le i \le n$, by grid sampling with step size s. Then, for each reference patch $I_{x_i^r}$, it searches in its neighborhood for up to m best matched patches such that each matched patch I_x satisfies $D(I_{x_i^r}, I_x) = \|\Psi_{2D}(I_{x_i^r}) -$

[2]In the case the patch number is even, where there are two mediate values at each pixel, the smaller-magnitude one is the optimal value.

$\Psi_{2D}(I_x)\|_2^2/d^2 \leq \epsilon$, where ϵ is a pre-defined constant. These matched patches form the i^{th} patch group I_{G_i}.

We set the threshold value ϵ such that the popular reference patches will have more (but up to m) matched patches than the rare ones. In this way, the grouping information will help to improve the recovery accuracy of the popular patches without harming that of the rare ones. Thus, we can increase the recovery accuracy of an image consisting of abundant correlated patches, by incorporating this non-local grouping information.

3.2. Non-Local Joint Recovery

Algorithm 1 Solve non-local joint recovery using inexact ALM-ADM

Require: C, S, B and $P_{ij}, \forall i, j$,
Ensure: I
1: $I^0 = g_1^0 = g_2^0 = b_1^0 = b_2^0 = R^0 = e^0 = \text{zeros}(M, N)$; $\overline{Z}_i^0 = \widehat{Z}_{ij}^0 = f_{ij} = \text{zeros}(N_1, N_1)$.
2: **while** I not converged **do**
3: Separate Estimate of Auxiliary Variables χ: $\chi^{k+1} \leftarrow \arg\min_\chi \mathcal{L}(I^k, \chi, \rho^k)$.
4: Joint Reconstruction of Image I: $(I^{k+1}) \leftarrow \arg\min \mathcal{L}(I, \chi^{k+1}, \rho^k)$.
5: Update of Lagragian Multipliers ρ: $b_l^{k+1} \leftarrow b_l^k - \tau\beta_1(g_l^{k+1} - D_l I^{k+1})$. $f_{ij}^{k+1} \leftarrow f_{ij}^k - \tau\beta_2(\overline{Z}_i^{k+1} + \widehat{Z}_{ij}^{k+1} - \Psi_{2D}(P_{ij}I^{k+1}))$. $e^{k+1} \leftarrow e^k - \tau\beta_3(R^{k+1} - CI^{k+1})$.
6: $k \leftarrow k + 1$.
7: **end while**

In this subsection, we present an efficient algorithm for the non-local recovery using non-local joint sparsity (NLJS). It is straightforward to extend it to solve the recovery problem using non-local wavelet sparsity (NLWS).

We denote the j^{th} patch in the i^{th} group as $I_{G_{ij}}$ and its extraction matrix as $P_{ij} \in \mathbb{R}^{d^2 \times MN}$, i.e., a binary matrix each row of which has one nonzero entry "1". Thus, we obtain $I_{G_{ij}} = P_{ij}I$, where I and $I_{G_{ij}}$ are vectorized, for simplicity of notation. For each patch $I_{G_{ij}}$, there is a common component $\overline{Z}_i$ and an innovation component $\widehat{Z}_{ij}$ such that $\overline{Z}_i + \widehat{Z}_{ij} = \Psi_{2D}(P_{ij}I)$. Accordingly, for the patch group I_{G_i}, the 3D cube of innovation components $\widehat{Z}_i = [\widehat{Z}_{i1} : ... : \widehat{Z}_{im_i}]$. Thus, the non-local recovery problem is formulated as follows:

$$\min_{I,\overline{Z}_i,\widehat{Z}_{ij}} \sum_{l=1}^{2} \|D_l I\|_1 + \lambda \sum_{i=1}^{n} (\|\overline{Z}_i\|_1 + \eta \sum_{j=1}^{m_i} \|\widehat{Z}_{ij}\|_1),$$
$$\text{s. t.} \quad SCI = B, \overline{Z}_i + \widehat{Z}_{ij} = \Psi_{2D}(P_{ij}I), \forall i, j. \quad (6)$$

It is very difficult to directly solve this constrained optimization problem consisting of multiple non-differentiable sparsity-inducing terms (ℓ_1 norm). So, we employ the augmented Lagrangian method-alternating direction method (ALM-ADM), also called the alternating direction method of multipliers (ADMM) in [3], to divide this complicated problem into simpler sub-problems and addresses them iteratively. This ALM-ADM algorithm has been widely used in compressive sensing [22]. By adding a set of auxiliary variables $\chi \triangleq \{g_1, g_2, \overline{Z}_i, \widehat{Z}_i, R\}$, the non-local recovery problem can be reformulated as

$$\min_{I,\chi} \sum_{l=1}^{2} \|g_l\|_1 + \lambda \sum_{i=1}^{n} (\|\overline{Z}_i\|_1 + \eta \sum_{j=1}^{m_i} \|\widehat{Z}_{ij}\|_1) \quad \text{s. t.}$$
$$\begin{cases} R = CI \\ SR = B \end{cases}, \; g_l = D_l I, \overline{Z}_i + \widehat{Z}_{ij} = \Psi_{2D}(P_{ij}I), \forall i, j, l. \quad (7)$$

This objective function given has the desirable property that it is separable in two groups of variables— the image I and its auxiliary variables $\chi = \{g_1, g_2, R, \overline{Z}_i, \widehat{Z}_i\}$. Thus, this function can be minimized over one group of variables by fixing the other group. Let $\rho \triangleq \{b_1, b_2, f_{ij}, e\}$ be a set of Lagrangian multipliers, we can write the Lagrangian function of this equality-constrained problem as follows.

$$\mathcal{L}(I, \chi, \rho) = \sum_{l=1}^{2} \|g_l\|_1 + \lambda \sum_{i=1}^{n} (\|\overline{Z}_i\|_1 + \eta \sum_{j=1}^{m_i} \|\widehat{Z}_{ij}\|_1)$$
$$+ \frac{\beta_2}{2} \sum_{i,j} \|\Psi_{2D}(P_{ij}I) - \overline{Z}_i - \widehat{Z}_{ij} + \frac{f_{ij}}{\beta_2}\|_{2,W_{ij}}^2$$
$$+ \frac{\beta_1}{2} \sum_{l=1}^{2} \|D_l I - g_l + \frac{b_l}{\beta_1}\|_2^2 + \frac{\beta_3}{2} \|CI - R + \frac{e}{\beta_3}\|_2^2, \quad (8)$$

where $SR = B$, β_1, β_2 and β_3 are large constant (e.g., 100), and P_{ij}^T is the transpose matrix of P_{ij}. We will explain the weighted ℓ_2 norm $\|\cdot\|_{2,W_{ij}}^2 = (\cdot)^T W_{ij}(\cdot)$ and its weight (diagonal) matrix $W_{ij} \in \mathbb{R}^{d^2 \times d^2}$ in the section Joint Reconstruction of Image I.

As shown in Algorithm 1, after just one round of alternatively minimizing the Lagrangian function $\mathcal{L}(I, \chi, \rho)$ with respect to I and χ, the multiplier ρ is updated immediately with a step length τ. Thus, the ALM-ADM algorithm has three iterative steps—(1) separate estimate of auxiliary variables χ, (2) joint reconstruction of image I and (3) update of Lagrangian multipliers ρ. With fixed Lagrangian multpliers ρ^k, we only update the image I and its auxiliary variables χ in one round of iteration. Therefore, this algorithm is also called inexact ALM-ADM [13, 22]. According to the theoretical analysis in [12], the inexact ALM-ADM is guaranteed to converge at $\tau \in (0, \frac{1+\sqrt{5}}{2})$, under certain technical assumptions.

3.3. Separate Estimate of Auxiliary Variables

In this subsection, we discuss how to estimate the auxiliary variables χ from a given image I by minimizing their

sparsity-inducing ℓ_1-norm. Given a typical ℓ_1-norm minimization problem $\min_a \|a\|_1 + \frac{\beta}{2}\|a-b\|_2^2$, it has a closed-form solution $\hat{a} = \mathbb{S}_{\frac{1}{\beta}}(b)$, where $\mathbb{S}_{\frac{1}{\beta}}(b)$ is the soft thresholding function defined as $\max\{\text{abs}(b) - \frac{1}{\beta}, 0\}. * \text{sign}(b)$. Accordingly, we respectively update the partial gradients g_1 and g_2, the common patch $\overline{Z}_i$ and the innovation component $\widehat{Z}_i$ as follows:

$$g_l^{k+1} = \mathbb{S}_{\frac{1}{\beta_1}}(D_l I^k + \frac{b_l^k}{\beta_1}). \quad (9)$$

$$\overline{Z}_i^{k+1} = \mathbb{S}_{\frac{1}{\beta_2}}(\frac{1}{m_i}\sum_{j=1}^{m_i}\Psi_{2D}(P_{ij}I^k) - \widehat{Z}_{ij}^k + \frac{f_{ij}^k}{\beta_2}). \quad (10)$$

$$\widehat{Z}_{ij}^{k+1} = \mathbb{S}_{\frac{1}{\beta_2}}(\Psi_{2D}(P_{ij}I^k) - \overline{Z}_i^{k+1} + \frac{f_{ij}^k}{\beta_2}). \quad (11)$$

Theoretically, to guarantee the convergence of the NLCS algorithm, we need to iteratively update the pair $(\overline{Z}_i, \widehat{Z}_i)$ until convergence. In practice, for computational efficiency, we can apply only one round of updating $(\overline{Z}_{i,i})$, which is shown to have stable convergence behavior in experimental results.

As for the circulant sample R, we first obtain it by applying circulant sampling on the image I^k and then set its subsamples at locations X (decided by B and its subsampling matrix S) as the measurements B.

$$R^{k+1} = CI^{k+1} - e^{k+1}/\beta_3. \quad (12)$$

$$R^{k+1}[X] = B. \quad (13)$$

3.4. Joint Reconstruction of Image I

In this section, we discuss the joint reconstruction of the image I from its auxiliary variables $\chi = \{g_1, g_2, R, \overline{Z}_i, \widehat{Z}_i, 1 \le i \le n\}$.

After the first step, we have a set of patch estimates $\Psi_{2D}(I_{G_{ij}}) = \overline{Z}_i + \widehat{Z}_{ij}$, which constitue an over-complete and spatially non-uniform representation of the image I. The conventional way to recover the image I is aggregating all the patch estimates using a weighted average. Following this aggregation method, we impose a weighted ℓ_2 norm regularization $\|(P_{ij}I) - \overline{Z}_i - \widehat{Z}_{ij} + \frac{f_{ij}}{\beta_2}\|_{2,W_{ij}}^2$ in Eq. (8), where $\|\cdot\|_{2,W_{ij}}^2 = (\cdot)^T W_{ij}(\cdot)$. The weight W_{ij} is defined to yield larger regularization weight on the rare patches and smaller weight on the popular ones. Specifically, for each pixel in the patch $I_{G_{ij}} = P_{ij}I$, we set its regularization weight as the inverse of its sampling frequency over all the patches $\{I_{G_{ij}}\}_{1\le i\le n, 1\le j\le m_i}$, i.e., $W_{ij} = P_{ij}(\sum_{i,j} P_{ij}^T P_{ij})^{-1} P_{ij}^T$.

By setting the derivative of the Lagrangian function $\mathcal{L}(I, \chi, \rho)$ with respect to I to be zero, we obtain the following condition on I.

$$\beta_1 \sum_{l=1}^{2} D_l^T(D_l I - g_l + \frac{b_l}{\beta_1}) + \beta_3 C^T(CI - R + \frac{e}{\beta_3}) + \beta_2(I - \sum_{i,j}\frac{P_{ij}^T \Psi_{2D}^{-1}(\overline{Z}_i + \widehat{Z}_{ij} - \frac{f_{ij}^k}{\beta_2})}{\sum_{i,j} P_{ij}^T P_{ij}}) = 0 \quad (14)$$

Since $D_1^T D_1, D_2^T D_2$ and $C^T C$ are circulant matrices and multiplying each of them with I is equivalent to some convolution on I. Thus, we can effieciently recover I^{k+1} by using Fast Fourier Transform to solve Eq. (14).

4. Experimental Results

Figure 2. Images (from top-left to bottom-right): Barbara, Cameraman, Lena, House, Building, TrainStation, Bone and Brain.

Our proposed NLCS focuses on compressive recovery, and thus it can take the measurements obtained from existing compressive imaging cameras, e.g., single-pixel camera [9]. In this work, instead of bothering with the imaging system, we simulate the compressive measurement and evaluate the performance of our NLCS algorithm, compared with the traditional 2DCS methods. Our NLCS explores the non-local sparsity of an image for compressive recovery, which has shown to be an intrinsic property of many modalities of image signals, such as visible light natural images [4, 16, 7] and medical images [11]. More generally, our NLCS is applicable to any other signals as long as the patch correlation assumption holds. To validate the effectiveness of our algorithm, similar to previous works [11, 10], we select six diverse natural images (Barbara, Cameraman, Lena, House, Building, and TrainStation) and two medical images (Bone and Brain) for evaluation[3], as shown in Figure 2.

For computational efficiency, we employ circulant sampling on these images and obtain the sampled data $B = SCI$ at different sampling rates. NLCS using NLWS, NLCS using NLJS and the previous 2DCS are respectively applied to recover the images from their sampled data B.

[3]Although it is desirable to evaluate the algorithm's performance on signals directly from medical imaging or satellite imaging systems, we are limited by such data access. Experiment results with visible light images should generalize to other modalities of signals with non-local sparsity.

Peak signal-to-noise ratio (PSNR) is used to gauge the recovery accuracy. In 2DCS (Eq. (1)), we empirically set the weight parameter for 2D wavelet sparsity as $\lambda = 0.66$ for all eight images. Then, our NLCS starts with the initialization image recovered by 2DCS at $\lambda = 0.66$ and iterates between non-local patch grouping and non-local joint recovery. In non-local grouping, we set the patch size $d = 8$, the step size $s = 4$ and the maximal group size $m = 8$. In non-local recovery, we empirically set the same parameters for both sparsity measures NLWS and NLJS: the weight $\lambda = 2$ in Eq. (2), the over-regularization parameter $\beta_1 = \beta_2 = \beta_3 = 100$ in Eq. (8), and $\tau = 1.618$ in Algorithm 1. To recover the groups of highly correlated patches for NLJS, we set the weight on the innovation component as $\eta = 10$ in Eq. (5).

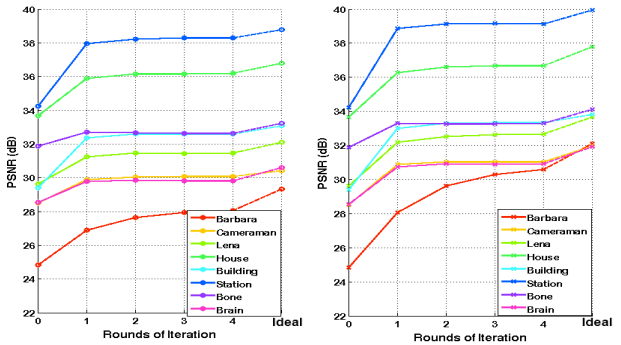

(a) (b)

Figure 3. Convergence behavior of the proposed NLCS—recovery accuracy vs. the iteration number (up to 4). (a) NLCS using NLWS and (b) NLCS using NLJS. We use the dotted line to connect the NLCS recovery at iteration 4 to the ideal case where the grouping information is extracted from the ground truth.

4.1. Quantative Evaluation

Admittedly, there is no theoretical guarantee that our NLCS algorithm can obtain global optimum by iterating between non-local grouping and non-local recovery. In Figure 3, we plot the curves of recovery accuracy of NLCS using NLWS and NLJS with respect to the number of iterations at sampling rate of 20% for the eight test images. Note that the recovery results at iteration zero are those from 2DCS. As shown, NLCS consistently improves the recovery accuracy and typically converges in about four iterations. To further evaluate the effectiveness of our NLCS, we also list the results of its *ideal* case where the grouping information is extracted exactly from the original image, which is the theoretical limit of our NLCS algorithm. After reaching its stable recovery accuracy, our NLCS using NLJS is only about 1 dB lower than the ideal cases, indicating that the NLCS is effective in recovering both the image and its non-local grouping information. On the images with many repeating structures (e.g., Barbara), NLCS has better performance and also takes more iterations to converge. On these eight images at sampling rate 10%, the final recovery accuracy of NLCS with NLWS is on average 2.56 dB higher than that of 2DCS, while the final accuracy of NLCS with NLJS is on average 3.80 dB higher than 2DCS. These significant improvement margins demonstrate the effectiveness of the non-local sparsity prior for image CS. Compared with NLWS, our NLJS is more effective as it models the patch group as a rank-1 subspace up to some sparse errors.

For complete comparison, NLCS with NLWS/NLJS and 2DCS are evaluated on the eight images at varying sampling rate ($10\%, 20\%$ and 30%), as shown in Table 1. NLCS using NLJS consistently outperforms NLCS using NLWS, and both significantly improves over 2DCS at all sampling rates. As expected, NLCS using NLJS achieve larger improvements in terms of recovery accuracy on images that contain sharp edges (e.g., Building and TrainStation), and images with repeating structures (e.g. Barbara). On Barbara, the improvement is even as large as 6.81 dB over 2DCS at sampling rate 30%. On TrainStation, NLCS using NLJS can improve the recovery accuracy up to 36.07 dB at the sampling rate of only 10%, suggesting that our new algorithm is much more practical compared with 2DCS.

4.2. Visual Quality Evaluation

To further compare our NLCS with 2DCS, we present their visual recovery results on some typical images. As shown in Figure 4 (note that images are better viewed in PDF), despite as the-state-of-the-art in image CS, 2DCS recovers the Barbara image with texture and edges (or boundaries) blurred. By adding the regularization on the non-local patch correlation, both NLCS using NLJS and NLWS can recover the Barbara image with correct texture and sharp edges. In the zoomed-in regions, we can see NLJS is superior to NLWS in recovering local details. The error maps of visual recovery also show that NLJS is better than NLWS and both significantly outperform 2DCS. In addition to the Barbara image that is full of repeating structures, we also show the visual recovery results on image TrainStation, classical image Cameraman and the medical image Brain in Figure 5. On these images, 2DCS always fails to recover details well, while the results from our NLCS hold much better fidelity to the ground truth and NLJS is always superior to NLWS as non-local sparsity prior.

5. Conclusion

In this paper, we have proposed a non-local compressive sampling (NLCS) recovery method that exploits non-local patch correlation and local piecewise smoothness in a neat optimization work. Our NLCS recovery method is shown to significantly reduce the required sampling rate of

Table 1. Evaluation of the proposed NLCS using NLWS, NLJS with 2DCS in terms of PNSR (dB) at varying sampling rate. We run NLCS using NLWS and NLJS in four rounds of iterations and compare them with 2DCS (the improvement over 2DCS is highlighted as "gain") and their ideal case where the grouping information is known.

Image	Sampling	2DCS	NLWS						NLJS					
			round1	round2	round3	round4	gain	ideal	round1	round2	round3	round4	gain	ideal
Barbara	10 %	22.73	23.78	24.11	24.24	24.29	**1.56**	25.85	24.28	25.07	25.48	25.70	**2.97**	28.60
	20 %	24.85	26.91	27.65	27.94	28.06	**3.21**	29.34	28.07	29.63	30.31	30.60	**5.75**	32.12
	30 %	27.01	30.17	31.12	31.37	31.43	**4.42**	32.23	31.69	33.35	33.76	33.83	**6.81**	34.35
Cameraman	10 %	24.70	25.97	26.26	26.30	26.31	**1.61**	27.16	26.89	27.44	27.62	27.67	**2.97**	29.01
	20 %	28.52	29.89	30.05	30.08	30.10	**1.58**	30.42	30.88	31.04	31.04	31.02	**2.50**	31.98
	30 %	31.46	32.65	32.75	32.81	32.80	**1.34**	32.99	33.61	33.78	33.75	33.73	**2.27**	34.48
Lena	10 %	26.24	27.55	27.81	27.80	27.85	**1.61**	28.68	28.37	28.80	29.90	28.92	**2.68**	30.40
	20 %	29.63	31.23	31.46	31.45	31.48	**1.85**	32.12	32.17	32.54	32.64	32.67	**3.03**	33.67
	30 %	32.42	34.23	34.38	34.40	34.38	**1.96**	34.85	35.09	35.37	35.40	35.41	**2.99**	36.12
House	10 %	30.39	32.84	33.19	33.30	33.31	**2.92**	34.08	33.74	34.24	34.33	34.37	**3.98**	35.20
	20 %	33.69	35.91	36.14	36.17	36.19	**2.50**	36.78	36.28	36.59	36.66	36.69	**3.00**	37.80
	30 %	35.81	37.88	38.18	38.24	38.26	**2.45**	38.82	38.18	38.64	38.83	38.92	**3.11**	39.86
Building	10 %	24.78	28.19	28.95	28.97	28.96	**4.18**	29.88	28.89	29.79	29.95	30.00	**5.22**	31.27
	20 %	29.39	32.37	32.59	32.61	32.60	**3.21**	33.08	33.00	33.31	33.35	33.35	**3.96**	33.83
	30 %	32.59	35.11	31.25	35.27	35.28	**2.69**	35.56	35.69	35.95	35.98	36.01	**3.42**	36.65
TrainStation	10 %	29.33	33.51	34.20	34.34	34.35	**5.02**	35.27	34.97	35.91	36.07	36.07	**6.74**	36.86
	20 %	34.24	37.96	38.23	38.30	38.31	**4.07**	38.77	38.86	39.14	39.15	39.14	**4.90**	39.94
	30 %	37.36	40.40	40.49	40.50	40.50	**3.14**	40.92	41.03	41.17	41.19	41.19	**3.83**	41.80
Bone	10 %	27.37	29.07	29.26	29.27	29.25	**1.88**	30.19	30.00	30.24	30.24	30.26	**2.89**	31.58
	20 %	31.88	32.71	32.68	32.65	32.63	**0.75**	33.24	33.31	33.27	33.28	33.30	**1.42**	34.10
	30 %	34.61	35.11	35.10	35.06	35.06	**0.45**	35.43	35.58	35.57	35.59	35.59	**0.98**	36.02
Brain	10 %	24.14	25.51	25.77	25.82	25.82	**1.68**	26.97	26.61	27.02	27.07	27.08	**2.94**	28.93
	20 %	28.54	29.77	29.84	29.83	29.83	**1.29**	30.60	30.75	30.91	30.90	30.91	**2.37**	31.94
	30 %	31.87	32.63	32.65	32.64	32.60	**0.73**	33.06	33.32	33.36	33.37	33.37	**1.50**	34.05

image CS and thus can greatly improve the existing compressive imaging cameras. To impose the patch correlation prior, we propose two non-local sparsity measures—non-local wavelet sparsity (NLWS) and non-local joint sparsity (NLJS). NLJS is consistently better than NLWS in recovering sharp edges and fine textures. An efficient algorithm consisting of two iterative steps, non-local grouping and non-local recovery, is developed to solve the NLCS recovery problem. Motivated by these promising experimental results, we will apply our NLCS to real compressive imaging systems (e.g., MRI system) in the next step.

References

[1] D. Baron, M. B. Wakin, M. F. Duarte, S. Sarvotham, and R. G. Baraniuk. Distributed compressed sensing. *Preprint*, 2005.

[2] S. Bengio, F. Pereira, Y. Singer, and D. Strelow. Group sparse coding. In *NIPS*, volume 22, pages 82–89, 2009.

[3] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1122, 2011.

[4] A. Buades, B. Coll, and J. M. Morel. A non-local algorithm for image denoising. In *CVPR*, 2005.

[5] E. Candes. An introduction to compressive sampling. pages 221–230, 2008. IEEE Signal Processing Magazine.

[6] E. Candes, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):article 11, 2011.

[7] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3d transform-domain collaborative filtering. *TIP*, 16(8):2080–2095, 2007.

[8] D. Donoho. Compressed sensing. *IEEE Transaction on Information Theory*, 2006.

[9] M. F. Duarte, M. A. Davenport, D. Takhar, and et al. Single-pixel imaging via compressive sampling. *IEEE Signal Processing Magazine*, 25(2):83–91, 2008.

[10] J. M. Duarte-Carvajalino and G. Sapiro. Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization. *TIP*, 2009.

[11] K. Egiazarian, A. Foi, and V. Katkovnik. Compressed sensing image reconstruction via recursive spatially adaptive filtering. In *ICIP*, 2007.

[12] R. Glowinski. *Numerical Methods for Nonlinear Variational Problems*. Springer-Verlag, 1984.

[13] Z. Lin, M. Chen, , and Y. Ma. The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices, 2010. UIUC Tech. Report UILU-ENG-09-2214.

[14] M. Lustig, D. Donoho, J. Santos, and J. Pauly. Compressed sensing MRI. *IEEE Sig. Proc. Magazine*, 2007.

[15] S. Ma, W. Yin, Y. Zhang, and A. Chakraborty. An efficient algorithm for compressed MR imaging using total variation and wavelets. In *CVPR*, 2008.

[16] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Non-local sparse models for image restoration. In *ICCV*, 2009.

[17] J. Romberg. Compressive sensing by random convolution. *SIAM Journal on Imaging Science*, 2009.

[18] L. Rudin and S. Osher. Total variation based image restoration with free local constraints. In *ICIP*, 1994.

[19] X. Shu and N. Ahuja. Hybrid compressive sampling via a new total variation TVL1. In *ECCV*, 2010.

[20] X. Shu and N. Ahuja. Imaging via three-dimensional compressive sampling (3DCS). In *ICCV*, 2011.

[21] M. Wakin, J. Laska, M. Duarte, D. Baron, S. Sarvotham, D. Takhar, K. Kelly, and R. Baraniuk. Compressive imaging for video representation and coding. *Picture Coding Symposium*, 2006.

[22] J. Yang and Y. Zhang. Alternating direction algorithms for l1-problems in compressive sensing. *University CAAM Tech. Report TR09-37*, 2009.

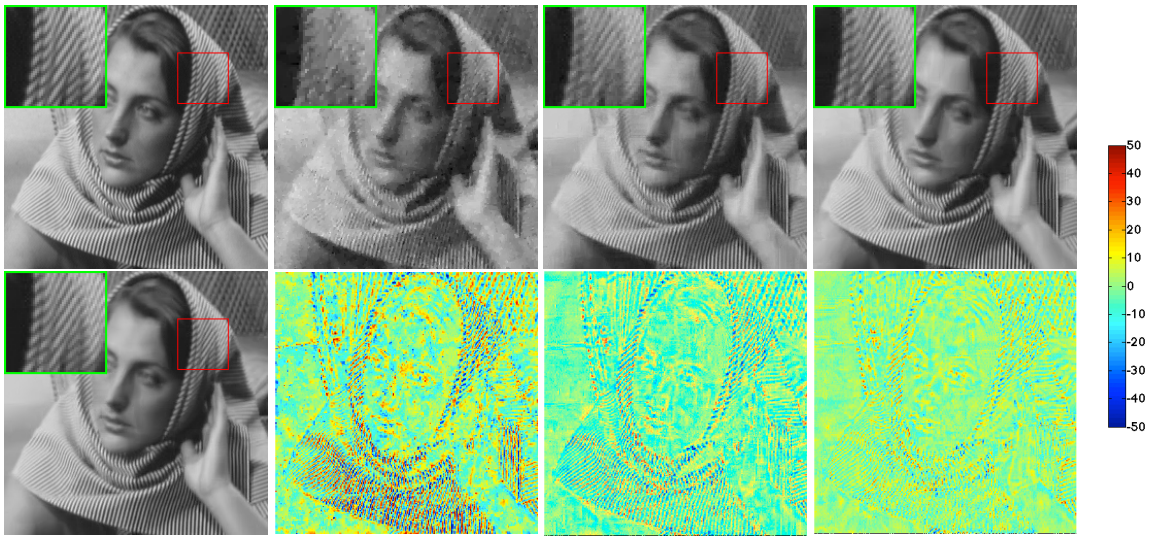

(a) (b) (c) (d)

Figure 4. Visual recovery of the proposed NLCS using NLWS and NLJS, in comparision with 2DCS on Barbara at 20% sampling rate. Top: (a) original image, recovered images by (b) 2DCS (PSNR:24.85), (c) NLWS (PSNR:28.06), and (d) NLJS (PSNR:30.60). Bottom: (a) image recovered by the ideal version of NLJS (PSNR:32.12), the error maps (with the color map on the left) of (b) 2DCS, (c) NLWS and (d) NLJS. **Note**: regions in green are zoomed-in versions of regions in red, and images are better viewed in PDF.

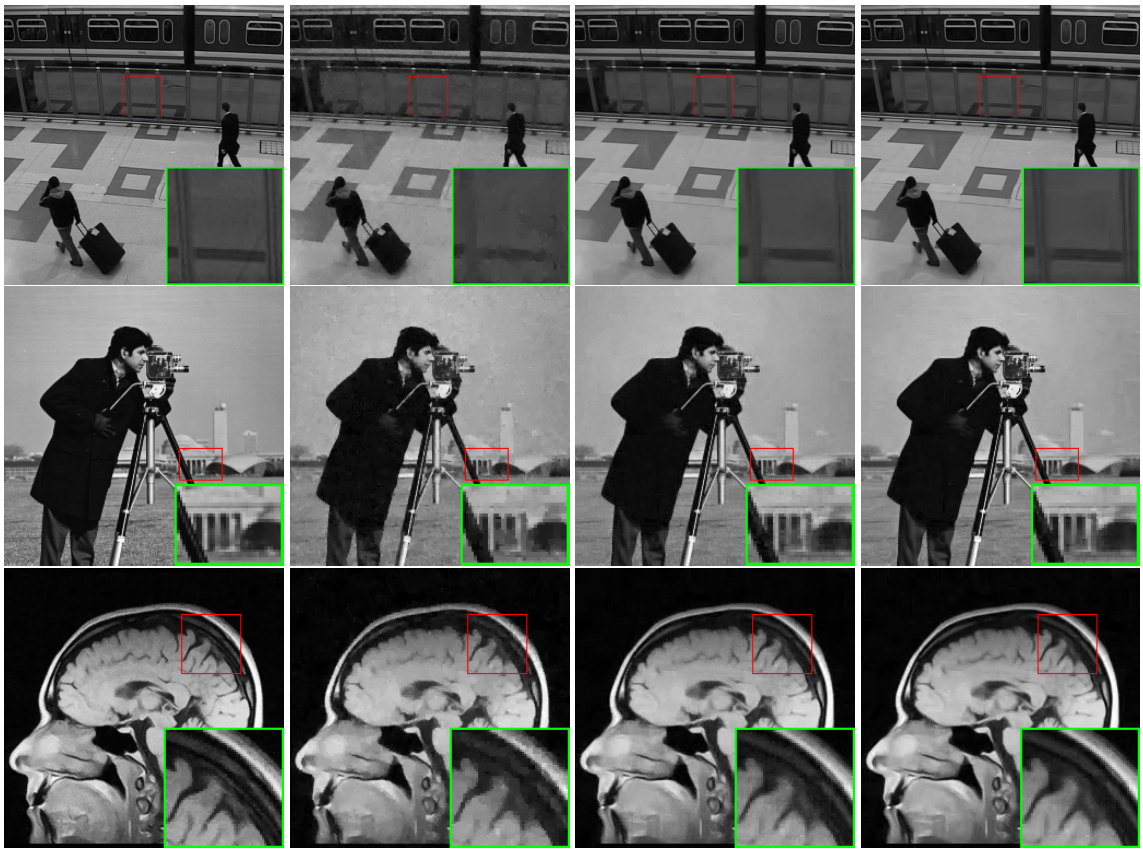

(a) (b) (c) (d)

Figure 5. Visual recovery of NLCS using NLWS and NLJS, in comparison with 2DCS on TrainStation (sampling rate: 10%), Cameraman (sampling rate: 20%) and Brain (sampling rate: 20%). (a) Original images, and images recovered by (b) 2DCS, (c) NLWS, and (d) NLJS. **Note**: regions in green are zoomed-in versions of regions in red, and images are better viewed in zoomed PDF.