Exploiting Non-local Low Rank Structure in Image Reconstruction

Evan Levine Tiffany Jou

***Abstract*—Constrained image models based on linear dependence are commonly used in high dimensional imaging and computer vision to exploit or extract structure, which can be expressed with low rank matrix approximations. Natural images also have a self-similarity property, where features tend to repeat themselves all over the image, and linear dependence relationships may be non-local. To exploit non-local linear dependence structure for image reconstruction, we develop a novel and flexible framework using low rank matrix reconstruction and a union of subspaces model, which is learned using subspace clustering. Subspace clustering makes weaker assumptions about the image than similar methods and also scales to very large-sized problems. We extend this approach for tensors, which are natural representations for images. We demonstrate a benefit of non-local low rank modeling for image denoising and reconstruction of MRI and light field images.**

***Index Terms*—image reconstruction, compressed sensing, MRI, light-field imaging**

I. Introduction

Many natural images have properties of self-similarity, where structures from image fragments tend to repeat themselves. Unlike methods based on models of local, transient phenomena such as wavelets, non-local methods use pixels from anywhere in the image to exploit internal correspondences. This type of non-local redundancy plays a role in many computer vision tasks such as object recognition, texture classification, tracking, alignment, and temporal motion segmentation [1]. In image processing, non-local image priors have achieved state-of-the-art performance in image denoising and compressed sensing. [2], [3], [4]

Many non-local methods such as BM3D [2] use two steps: 1) grouping of similar image blocks and 2) collaborative filtering within groups to estimate the true image. The aim is to group similar blocks to construct arrays that are sparsely representable, and collaborative filtering exploits this sparsity by transform domain shrinkage. Grouping typically relies on a distance measure and collaborative filtering on a 3D transform (e.g. discrete cosine transform). Clearly, these steps are highly dependent, but there is no general answer as to how they could be systematically integrated. In practice, grouping strategies based on Euclidean distance are commonly used. One non-local method similar in spirit to BM3D was recently proposed for compressed sensing image recovery and shown to significantly outperform competing approaches. The method uses patch grouping based on Euclidean distance and collaborative filtering with singular value thresholding.

We consider reconstructing an image X from additive white Gaussian noise-corrupted measurements Y observed through a linear measurement model A. We also consider a related denoising problem corresponding to $A = I$. A non-local method decomposes X into a collection of vectorized patches of equal size $\{x_j \in \mathbb{R}^D\}_{j=1}^N$ that may overlap. Let $\{R_c\}_{c=1}^n$ be a set of linear operators such that R_c extracts patches $x_{i_{1,c}}, x_{i_{1,c}}, \cdots, x_{i_{d_c,c}}$ and reformats them into a matrix $R_c X = \begin{bmatrix} x_{i_{1,c}} & x_{i_{1,c}} & \cdots & x_{i_{d_c,c}} \end{bmatrix}$. Non-local methods find patch clustering operators R_c and impose a set of penalties $J_c(\cdot)$ on each patch cluster, solving a problem of the form

$$P0 : \hat{X} = \arg\min_X \|Y - AX\|_F^2 + \sum_c J_c(R_c X) \quad (1)$$

In the denoising setting, transform domain collaborative filtering of each patch cluster (as in BM3D) can be represented by sparsifying transform matrices Ψ_1 and Ψ_2 and $J_c(R) = \lambda_c \|\Psi_1 R \Psi_2^T\|_1$. In non-local low rank methods, J_c is a low-rank inducing penalty.

Prior work in non-local low rank regularization approximates $R_c X$ with low rank matrices $\{L_c\}$ by attempting to solve P1 for X and $\{L_c\}$. Rather than jointly minimizing over $\{R_c\}$, a patch group is defined for each image patch based on all patches within a fixed Euclidean distance, and alternating minimization is applied to the variables $\{L_c\}$ and X.

$$P1 : \hat{X} = \arg\min_X \|Y - AX\|_F^2 + \quad (2)$$

$$\eta \{ \sum_c \|R_c X - L_c\|_F^2 + \lambda \|L_c\|_* \} \quad (3)$$

Local low rank regularization has been applied in dynamic MRI to exploit temporal redundancy [5], [6], [7], [8]. In this setting, the definition of a patch is a signal intensity vs. time curve from a single voxel. Time curves are grouped using fixed non-overlapping neighborhoods of voxels (e.g. of size 6×6 voxels) that uniformly tile the image domain. Unlike P1, the formulation P2 follows P0 with $J_c(R) = \|R\|*$ and does not use L_c

$$P2 : \arg\min_X \|Y - AX\|_F^2 + \sum_c \lambda_c \|R_c X\|_* \quad (4)$$

In this work, we propose a novel reconstruction model that exploits non-local low rank structure. Much like other methods using low rank regularization, patch groups are sequentially segmented and approximated with low rank matrices. However, we integrate patch grouping and low rank approximation using a subspace clustering approach, thus removing the

distance parameter previously required for patch grouping. Following patch grouping via subspace clustering, images are reconstructed by solving P2 using an alternating direction method of multipliers (ADMM) algorithm. Numerical experiments show excellent reconstruction accuracy in grayscale image denoising, compressed sensing MRI, and light field imaging.

We list four novel contributions of this work.

1) *Distance-free non-local low rank via subspace clustering.* We introduce a new model to integrate patch clustering and the non-local low rank prior in the framework of subspace clustering.
2) *Extension of non-local low rank priors to tensors.* We extend the non-local low rank framework to reconstruct high-dimensional images by keeping tensors intact instead of reformatting to 2D matrices.
3) *Application of non-local low rank priors to MRI.* Dynamic MRI reconstruction with local low rank priors has gained significant attention. To our knowledge, non-local low rank regularization has not been applied to multidimensional MRI.
4) *Application of non-local low rank tensor reconstruction to Light Field Imaging.* We apply the non-local low rank tensor reconstruction method to recovering missing samples in 4-D Light Field images.

II. Theory

We briefly analyze the non-local low rank image model and show a connection to subspace clustering. Because patch clusters are low rank, we can approximate them using "tall" matrices U_c and V_c, writing

$$R_c X \approx U_c V_c^T \tag{5}$$

$$\begin{bmatrix} | & | & & | \\ x_{i_{1,c}} & x_{i_{1,c}} & \cdots & x_{i_{d_c,c}} \\ | & | & & | \end{bmatrix} \approx U_c V_c^T \tag{6}$$

This can be rewritten

$$x_{i_{j,c}} \approx U_c (V_c^T)_{:,i_{j,c}}, j = 1, 2, ..., d_c \tag{7}$$

We interepret this as follows. Non-local low rank regularization approximates each patch with a low-dimensional representation $(V_c^T)_{:,i_{j,c}}$ and a basis U_c, which depend on the cluster to which it belongs. Patches in cluster c are therefore constrained to lie in a subspace S_c of dimension $d_c = \dim(S_c)$ parameterized by U_c as

$$S_c = \{x \in \mathbb{R}^D : x = U_c v\} \tag{8}$$

.

for some $v \in \mathbb{R}^{d_c}$. We can see that the low rank approximations (7) correspond exactly to the subspaces (8), illustrating that the union of subspaces $\cup_{c=1}^{n} S_c$ is the image representation from non-local low rank regularization. The problem is to recover both a union of subspaces and a segmentation.

Prior work has used Euclidean distance-based block-matching to cluster images patches and constrain clusters to lie in a subspace. [4] The procedure can be interpreted geometrically as growing an ℓ_2 ball around a data point (patch) and projecting the contained points to a low dimensional subspace. However, an ℓ_2 ball is not guaranteed to contain the underlying subspace in general, which suggests that it fails when patches are spread out over the subspace. Rather than developing problem-dependent distance measures for the same two-step procedure, we propose to side-step the choice of an explicit distance measure and use only the union of subspaces model, effectively integrating the two steps. Assuming only a union of subspaces allows arbitrary distributions of the patches within the subspaces. This problem is precisely what subspace clustering attempts to solve, which has been investigated recently in computer vision to estimate data from a union of subspaces [9].

We illustrate the subspace clustering approach in Figure 1. The basic idea is to use subspace clustering to segment the patches and low rank regularization cluster-wise to project them onto the original subspaces. In this example, the image consists of signal-intensity vs. time curves, which have been estimated previously using local low rank methods.

A. Subspace clustering

Many methods for subspace clustering have been developed and have been classified as algebraic, iterative, statistical, and spectral-clustering based [10]. Many methods are not robust to noise and outliers, are computationally expensive, face numerical issues, or are very heuristical. Sparse subspace clustering (SSC) is one spectral-clustering based method that is robust to noise and outliers and can achieve very accurate subspace segmentations [11]. Instead of clustering points based on angular or Euclidean distance, SSC invokes a sparsity criterion. Each data point X_i is approximated as a sparse linear (or affine) combination of all other data points with coefficients β, computed by solving the ℓ_1 minimization problem

$$\|X - X\beta\|_F^2 + \lambda\|\beta\|_1 \tag{9}$$

$$\text{subject to } \beta_{ii} = 0, i = 1, 2, ..., N \tag{10}$$

SSC produces a sparse affinity matrix $A_{ij} = |\beta_{ij}| + |\beta_{ji}|$ with nonzero elements corresponding to patches in the same subspace, and under certain conditions, it can be guaranteed that $A_{ij} = 0$ if patches i and j are in different subspaces [12]. The final subspace segmentation is obtained by applying spectral clustering [13] to A. Although SSC performs well in practice, its high computational complexity ($O(N^2)$) is a major disadvantage. In simple experiments with SSC, we found that typically all data points are expressed using only a few data points that are large in magnitude because points with high signal tend to encode the subspaces well. Interestingly, this redundancy in the solution suggests that the SSC solution is simpler than what 9 would lead us to believe and that more efficient methods are possible.

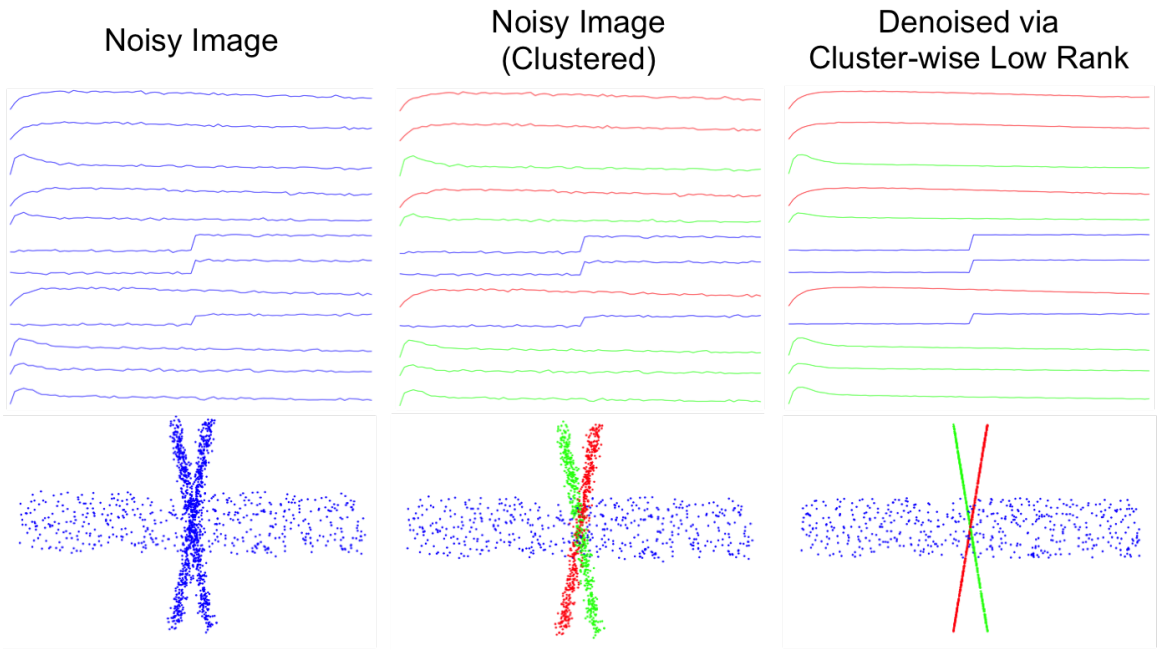

Fig. 1. In this illustration, an image consists of signal intensity-time curves for each pixel (top row), which can be represented as data points near a union of subspaces (bottom row). The proposed image model is based on a union of subspaces, which can be estimated by learning the subspaces via subspace clustering and cluster-wise low rank approximation.

Low-rank representation (LRR) is a similar method that extends SSC by replacing the ℓ_1 norm with a nuclear norm and can be more accurate. The coefficient matrix is computed from the following nuclear norm minimization

$$\|X - X\beta\|_F^2 + \lambda\|\beta\|_* \tag{11}$$

LRR also has high ($O(N^3)$) computational complexity. However, a three-step procedure with computational complexity $N \log N$ was recently proposed for fast low rank subspace segmentation (Fast LRSS). [14] Letting the rank-r SVD of X be $U_r\Sigma_r V_r^H$, the solution to (11) is $\beta^* = V_r S_\lambda(\Sigma_r)V_r^H$, where $\lambda = \frac{1}{\sigma_{r+1}^2}$ and

$$S_\lambda(x) = \max(0, 1 - \frac{1}{\lambda t^2}) \tag{12}$$

The solution can be computed efficiently with a partial SVD ($O(drN)$ complexity). To achieve $O(N \log N)$ complexity, Fast LRSS uses a sparse version of the affinity matrix defined by $A_{ij} = \frac{|\langle y_i, y_j\rangle|}{\|y_i\|_2\|y_j\|_2}$, where

$$y_i = (V_r\sqrt{S_\lambda(\Sigma)})_{i,:}. \tag{13}$$

B. Extension to Tensor Representation

In multidimensional imaging, images or image patches are naturally represented as tensors, and extending low rank methods to tensors can be challenging. Most existing approaches convert tensor problems to matrix problems, where images are reformatted (unfolded) as matrices, which are amenable to efficient and well-understood matrix analysis. The unfolding of a tensor along dimension i is defined as the matrix whose columns are formed as vectorized slices along dimension i. For example, 4D lightfield images can be unfolded by mapping slices along the spatial dimension Y to rows or similarly with slices along the viewing angle (Figure 2). Much previous work in tensor low rank modeling is based on enforcing low rank on different unfoldings of tensors. Often, one or multiple unfoldings may be used, and when multiple unfoldings are used, a sum of the nuclear norms of unfoldings is used, and this penalty can be viewed as a proxy for the tensor rank [15].

However, analyzing tensors with matrix-based methods has several limitations. If only a particular unfolding in dimension i is used, low rank in other dimensions are not exploited [16]. On the other hand, using multiple unfoldings and weighting them is fundamentally challenging and problem dependent. Since there can be multiple ways to reformat the data, often, it is unclear which unfoldings yield low rank matrices and which should participate in the reconstruction. Therefore, instead of exploiting low-rank in a reformatted matrix, the high-dimensional array should ideally be kept in its original form so that nonlocal redundancy could be exploited over all dimensions.

To our knowledge, current tensor completion strategies do not apply non-local clustering to low tensor rank models. In our project, we want to extend our non-local low rank model to deal with high-dimensional tensors in their original tensor form and explore the possibility of applying this model to high-dimensional imaging modalities such as Lightfield Imaging [17], [18]. For example, if we had a N-dimensional dataset, we cluster similar N-dimensional tensor patches of a

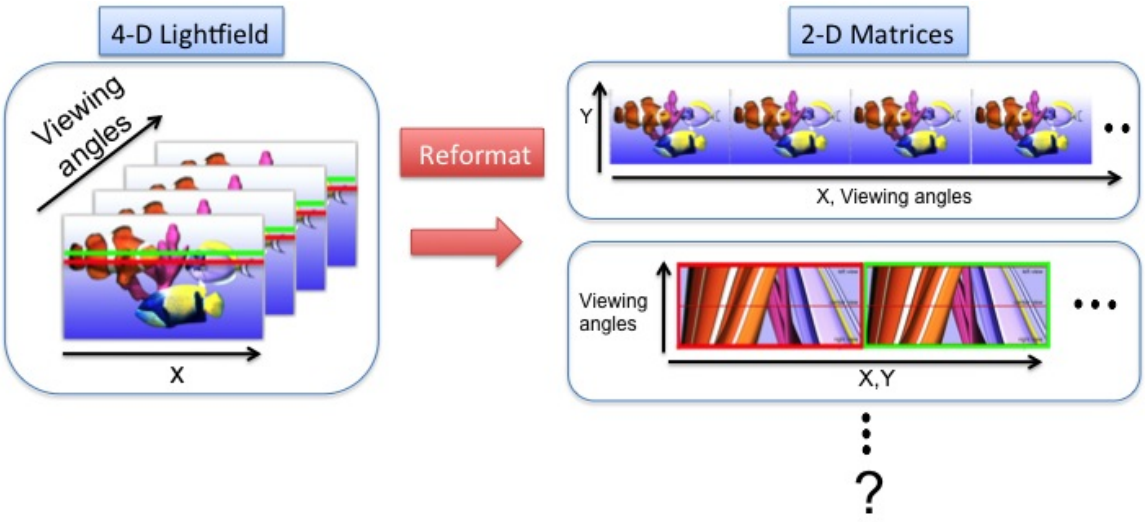

Fig. 2. There exists many different ways to reformat a 4-D lightfield into a 2-D matrix. Images adapted from EE367 course slides.

particular size among all possible tensor patches non-locally, concatenate them in a (N+1)-dimensional tensor, and enforce the low rank property on these clusters.

One important step in the matrix low-rank constraint reconstruction is the soft-thresholding step where singular values of a matrix are taken into account and thresholded to achieve a lower rank matrix. First of all, tensor rank is defined as the smallest number of rank-one tensors that the original tensor can be decomposed to and hold in equality [19]. This definition is analogous to the matrix rank definition. However, because computing the rank of a tensor is an NP hard problem, there is no expression for the convex envelope of tensor rank [20]. In addition, there is no high-dimensional tensor equivalent of Singular Value Decomposition (SVD). Therefore instead, we considered tensor CANDECOMP/PARAFAC-decomposition (CP-decomposition), which factorizes a tensor into a weighted sum of normalized rank-1 tensors [19]. There are other types of tensor decompositions, including Tucker-decomposition, which can be used to compute High-Order SVD (HOSVD). However, these methods are more associated with different dimension-*i* unfoldings. In our case, in order to consider the tensor block as a whole and to align with the tensor rank definition, we developed a method using the CP-decomposition.

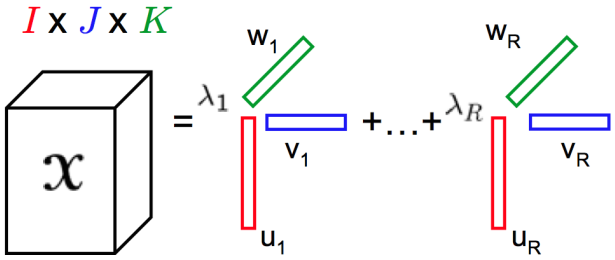

Fig. 3. CP-decomposition of 3-D tensor [21]

For a tensor with 4 dimensions, the rank R CP-decomposition can be written as:

$$\chi \approx \sum_{r=1}^{R} \lambda_r a_r \circ b_r \circ c_r \circ d_r \tag{14}$$

where $\circ$ represents vector outer product. The columns of the A,B,C matrices are normalized and a scaling vector λ is used to represent the weights, or contributions, of the different rank-1 tensor components.

The CP-decomposition can be used to compute a rank-r tensor estimate of the original tensor. However, as mentioned before, computing the rank of the original tensor is an NP-hard problem. Therefore, it is difficult to threshold the rank of the lower-rank tensor according to the original rank. Also, the properties of matrix SVD do no translate directly to CP-decomposition; for example, the columns of A, B, and C matrices of the CP-decomposition are not orthonormal and can be linearly dependent [21]. Therefore, for this project we developed a nonconvex heuristical algorithm by taking into account the λ vector of the CP-decomposition. Although the CP-decomposition λ weights are not equivalent to singular values, they represent the importance of the rank-1 tensor components and their contribution to the final decomposition sum. In our approach, we used the λ vector to threshold the 5-D tensor to be low rank.

C. Algorithms

To solve the unconstrained problem (4), we use variable splitting to transform it to an equivalent constrained problem, which can be cast in an augmented Lagrangian framework and efficiently solved using alternating minimization via the alternating direction method of multipliers (ADMM) [22]. Using the variable split $Z_c = R_c x$, the augmented Lagrangian is

$$L_\rho(X, Z, U) = \|Y - AX\|_F^2 + \sum_{R_c \in C} \lambda_c \|Z_c\|_* \quad (15)$$

$$+U_c^T(R_c X - Z_c) \quad (16)$$

$$+\rho/2\|R_c X - Z_c\|_F^2 \quad (17)$$

Applying ADMM yields Algorithm 1, which requires only functions that implement the forward operators A and R_c as well as their adjoints. The term $\sum_{R_c \in C} R_c^H R_c$ is a diagonal matrix with an image containing the number of times each pixel appears in a patch on the diagonal. In the special case where A is an undersampled Fourier matrix, we note that the Gram matrix ($A^H A$) is circulant and therefore diagonalizable in the Fourier basis, which allows the X-update to be implemented using a single forward and inverse multidimensional FFT and pixel-wise multiplication.

Algorithm 1 ADMM - P1

while stopping criteria false **do**
  $X^{k+1} \leftarrow (A^H A + \rho \sum_{R_c \in C} R_c^H R_c)^{-1}(A^H Y + \rho \sum_{R_c \in C} R_c^H (Z_c^k - U_c^k))$
  for $c \in \{1..K\}$ **do**
    $Z_c^{k+1} \leftarrow SVT(R_c X^{k+1} + U_c^k; \lambda_c/\rho)$
    $U_c^{k+1} \leftarrow U_c^k + \rho(R_c X^{k+1} - Z_c^{k+1})$
  end for
  $k \leftarrow k + 1$
end while
return x^k

III. Experiments and Results

We performed experiments with grayscale image denoising, compressed sensing in MRI, and light field image inpainting to compare different low rank priors. In denoising experiments, 8×8 image patches were used and overlapped at strides of 2 pixels ($4\times$ overlap). In experiments with MRI reconstruction, signal intensity curves, each from a single voxel, were treated as patches.

Local low rank priors currently used in MRI were compared using different MRI images. Non-local low rank priors were implemented by clustering patches from an initial image estimate (the original noisy image or minimum norm solution) and solving 4. The first non-local method was implemented using spectral clustering with Euclidean distance, and we refer to this method as ℓ_2**-NLLR**. A second non-local low rank method was implemented using subspace clustering with Fast LRSS to cluster patches, and we refer to this method as **Fast LRSS-NLLR**.

In experiments with grayscale denoising and MRI data, we selected λ_c automatically using methods based on Stein's unbiased risk estimator (SURE) [23], which requires the image noise to be i.i.d. Gaussian with a known standard deviation. We simply provided SURE with the measurement noise and verified that it chose nearly optimal regularization parameter selections in several of our experiments.

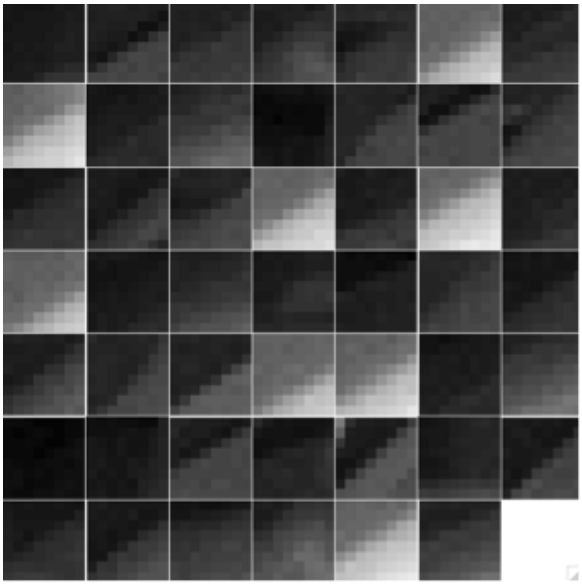

Fig. 4. Patch cluster from Fast LRSS-based subspace clustering showing intensity variation

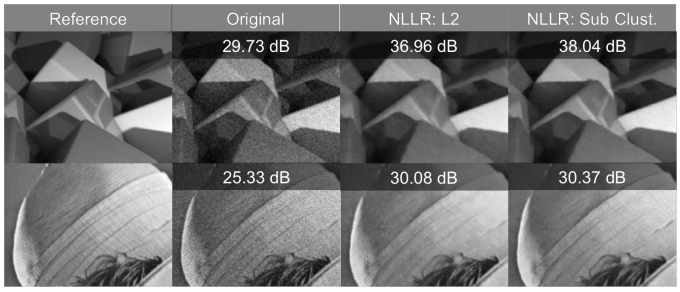

Fig. 5. Cubes and Lenna images denoised with non-local low rank methods

A. Experiments with noisy images

Numerical denoising experiments were performed using set of standard test images, another image "Cubes" consisting of cubes with different intensities, and a complex-valued brain image. Additive white Gaussian noise was added to each reference image to simulate SNR levels of 15, 17, 20, and 25 dB, and images were denoised with Fast LRSS-NLLR and ℓ_2-NLLR methods.

The example patch cluster obtained from Fast LRSS-based subspace clustering in Figure III-A shows that Fast LRSS can recover patches that lie in the same subspace but are not close in ℓ_2 distance. Since these patches are nearly the same diagonally oriented edge, only with varying intensities, this patch cluster admits a rank 1 (or 2 if there is a nonzero mean) approximation. However, these patches are not close in ℓ_2 distance due to the intensity variation.

Images denoised with Fast LRSS-NLLR achieved higher PSNR than images denoised with ℓ_2-NLLR in all cases except for the cameraman image. In the "Cubes" image, the difference (≈ 1.2 dB) was much more significant.

Image	Method	SNR (dB)			
		15	17	20	25
Lenna	ℓ_2-NLLR	28.66	30.08	32.16	35.16
	Fast LRSS-NLLR	28.90	30.37	32.45	35.56
Monarch	ℓ_2-NLLR	27.22	28.51	30.38	33.85
	Fast LRSS-NLLR	27.57	28.90	30.94	34.53
Parrots	ℓ_2-NLLR	28.04	29.33	31.49	34.58
	Fast LRSS-NLLR	28.13	29.59	31.77	35.26
Boats	ℓ_2-NLLR	26.47	27.70	29.58	32.85
	Fast LRSS-NLLR	26.88	28.10	30.01	33.33
Barbara	ℓ_2-NLLR	26.44	27.73	29.66	33.03
	Fast LRSS-NLLR	26.81	28.16	30.13	33.73
Foreman	ℓ_2-NLLR	27.63	29.73	32.03	34.96
	Fast LRSS-NLLR	28.24	29.72	32.06	35.17
House	ℓ_2-NLLR	27.69	29.91	31.84	35.08
	Fast LRSS-NLLR	28.84	30.42	32.36	35.40
Cameraman	ℓ_2-NLLR	27.28	28.62	30.63	34.13
	Fast LRSS-NLLR	27.36	28.60	30.48	33.95
Cubes	ℓ_2-NLLR	33.54	34.82	37.02	40.34
	Fast LRSS-NLLR	34.16	35.78	38.24	41.67
Brain	ℓ_2-NLLR	29.42	30.40	33.16	38.05
	Fast LRSS-NLLR	29.34	30.33	33.27	38.08

TABLE I
PSNR OF IMAGES DENOISED WITH NON-LOCAL METHODS SHOW THAT FAST LRSS-NLLR CAN IMPROVE UPON THE ℓ_2-BASED METHOD

B. Experiments with MRI data

Evaluating constrained medical imaging methods *in vivo* is challenging as it is often hard to obtain ground truth images or create realistic phantom images. One way to address this challenge is to use a variable flip angle (VFA) breast acquisition, which offers *in vivo* data with a ground truth. [24] In this approach, a static object is imaged as a flip angle acquisition parameter is varied, and because the signal in different tissues has a different variation with flip angle, this acquisition can be used to mimic spatial and temporal variations seen in dynamic imaging. VFA breast data was undersampled in flip angle - Fourier space using a variable density random undersampling pattern and reconstructed with local and non-local low rank methods, and results were compared to fully-sampled original images. Figure III-B shows images reconstructed with local and non-local low rank methods. Local low rank images show blocking artifacts (arrows), which obscure edges visible in either of the NLLR images.

Additional reconstruction were performed on data from a prospectively undersampled dynamic contrast-enhanced MRI acquisition as part of a liver protocol using a randomized sampling pattern [25] with 6% sampling. The system matrix A consisting of Fourier transform and receive channel sensitivities was estimated using the ESPIRiT calibration method [26], and a least squares solution was compared to the same method with additional penalties for spatial ℓ_1-Wavelet, local low rank, and non-local low rank. The reconstructed images in III-B show significant differences in image quality, and only NLLR images clearly depict a small vessel shown in 7b. Spatial ℓ_1 wavelet cannot recover a good image from data within the time frame, so some temporal regression with either LLR or NLLR is certainly necessary.

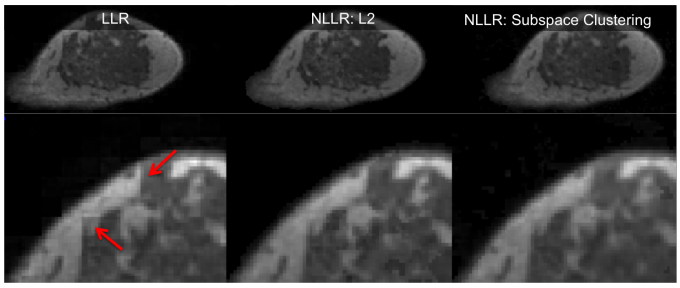

Fig. 6.

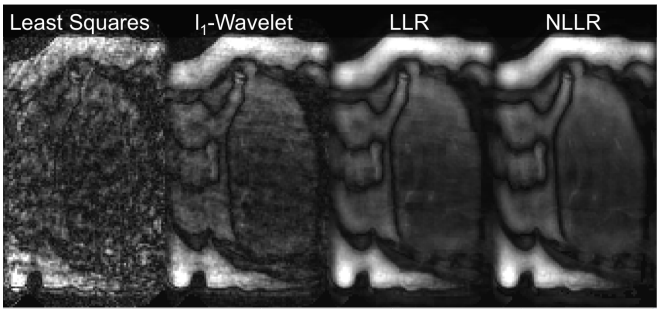

(a) Sagittal slices show DCE images reconstructed with four methods: a least squares solution, spatial ℓ_1-Wavelet regularization, local low rank (LLR) and ℓ_2-non-local low rank (NLLR).

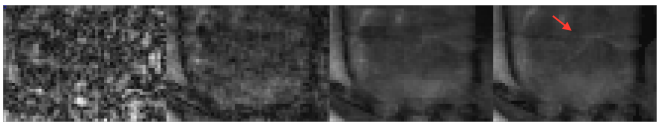

(b) ROIs from the images in (a) showing a small vessel visible in NLLR images

Fig. 7. Prospectively acquired DCE images

C. Experiments with Inpainting of Light field Images

We applied the nonlocal low rank tensor framework to reconstructing light field images with missing samples. To cluster the 4-D tensor patches, we vectorized the patches and clustered according to ℓ_2 norm distance. After clustering, the tensor patches were reformatted back to the original 4-D shape and similar patches were stacked in the 5th dimension to form a low rank tensor. For non-locally low-rank clustering, ideally we would like to find all the tensor patches in the tensor space that is within an ℓ_2 norm distance of the exemplar patch. However, this would increase computation time in the ADMM iteration loop. Therefore instead, we found the K nearest neighbors in terms of ℓ_2 norm of the tensor values.

To enforce low rank, we computed a rank-R CP-decomposition of the 5-D tensor cluster. R is an estimate of the desired rank of the resulting low-rank tensor. In these experiments, we used R = 20 as the starting parameter for the algorithm. A rank-20 CP-decomposition is computed, and the λ weights are sorted in descending order. Since λ values do not have the same charactersitcs as singular values, soft-thresholding these values do not yield good results. Instead, the number of λ weights needed to compose 85% of the

decomposition is noted as r. However, we cannot just take the first r components of the original rank-R decomposition as the best rank-r approximation, because the components of the best rank-r model must be found simultaneously [Reference: Tensor and applications]. Therefore, a new rank-r CP-decomposition is computed as the low rank tensor cluster to result in a better rank-r approximation.

In Figure 8 and Figure 9, we have the original light field image of a certain view, and the same image with 25% samples missing in the top rows. The left figures on the bottom rows were reconstructed by reformatting to a matrix before enforcing low rank. Results from these experiments will be referred to as "matrix NLLR". For these, clustering was done with (10,10,1,1) tensor patches. After clustering, the (10,10,6) matrices were then reformatted into (100,6) matrices before soft-thresholding the singular values. The right figures in the bottom row were reconstructed by keeping the tensor patches intact. These results will be referred to as "tensor NLLR". The original 4-D light field data of size (200,300,5,5) was split into 4-D tensor blocks of size (10,10,2,2). These tensor patches were then clustered with 6 other patches into 5-D tensors of size (10,10,2,2,6). Low tensor rank was then imposed on these clusters with the method described previously.

The matrix NLLR results are sharper and recover high frequency detail better, though the tensor NLLR results recover many of the edges and detailed features as well, even without taking into account which dimension is known to be redundant. The root-mean-square (RMS) errors comparing the matrix NLLR and the tensor NLLR reconstructions are included in the figures. For both of our data sets, our tensor NLLR reconstructions resulted in lower RMS errors. This can be seen by the many dotted patches in the reformatted matrix NLLR results, which may be due to the fact that some patches appear in more clusters than others. This didn't seem to be a problem with the tensor NLLR reconstruction, although both methods used the same number of nearest neighbors. Also, in Figure 9 and Figure 10, the matrix NLLR reconstructions appear to have more patchy artifacts and patch flickering between different views, which is not as apparent in the tensor NLLR reconstructions. Finally, the parallax information contained in the light field between different views is important to preserve. In our tensor NLLR reconstruction results, we did not see much loss in that information, as seen in Figure 10.

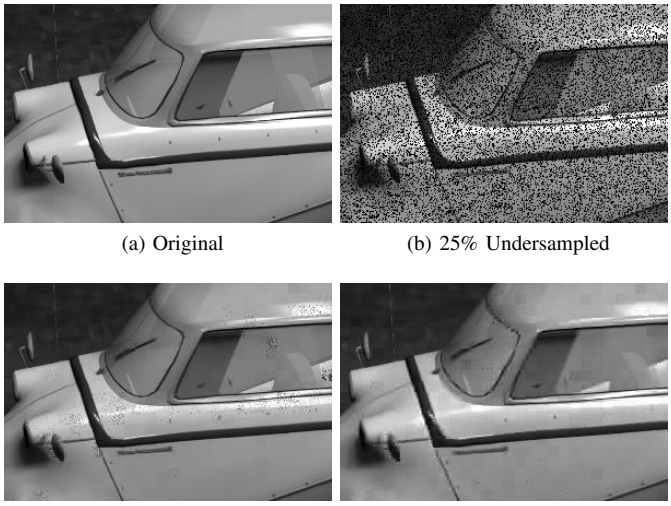

(a) Original (b) 25% Undersampled

(c) Matrix NLLR, RMS error=48.70 (d) Tensor NLLR, RMS error=39.13

Fig. 8. Inpainting of Lightfield Images

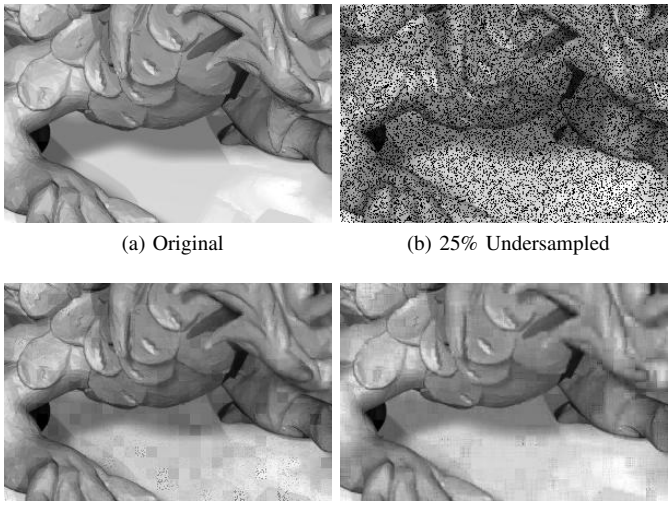

(a) Original (b) 25% Undersampled

(c) Matrix NLLR, RMS error=99.52 (d) Tensor NLLR, RMS error=55.19

Fig. 9. Inpainting of Lightfield Images

IV. Discussion

A. Comparison of ℓ_2-NLLR vs. Fast LRSS-NLLR

Denoising results suggest that patch clustering with ℓ_2-distance criteria can be improved with subspace clustering via Fast LRSS. The improvement can be explained by the failure of the ℓ_2 ball to include all points in the subspace, which results in failure to learn the subspaces accurately and a higher rank. Fast LRSS-NLLR uses only a union of subspaces model, which allows arbitrary distributions of patches within the subspaces. Fast LRSS allows patch clusters like the one shown in Figure III-A, which consists of edges with a variety intensities. Because the image has smooth spatial variations in intensity, patches are not close in ℓ_2 distance, violating the assumption of ℓ_2-NLLR. Fast LRSS performs significantly better than ℓ_2 because it is robust to intensity variations, while the improvement is less significant for the other images, which do not have significant intensity variations. Fast-LRSS may be beneficial for images with smooth spatial variations in intensity due to lighting or color. There are many sources of spatial variation in MRI such as field inhomogeneity, image phase variation, image contrast mechanisms, and physical properties of tissues.

Generalizing non-local low rank regularization via subspace clustering helps reduce problem-dependent parameter tuning such as the distance criterion, and one should prefer to use a more general method if the reconstruction is no worse. We did not find that Fast LRSS-NLLR clearly outperformed ℓ_2-NLLR in experiments with MRI data, which suggests that time curves

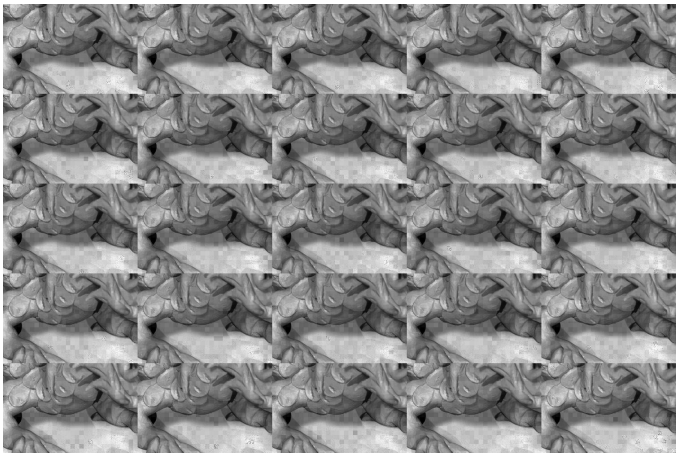

(a) Matrix NLLR

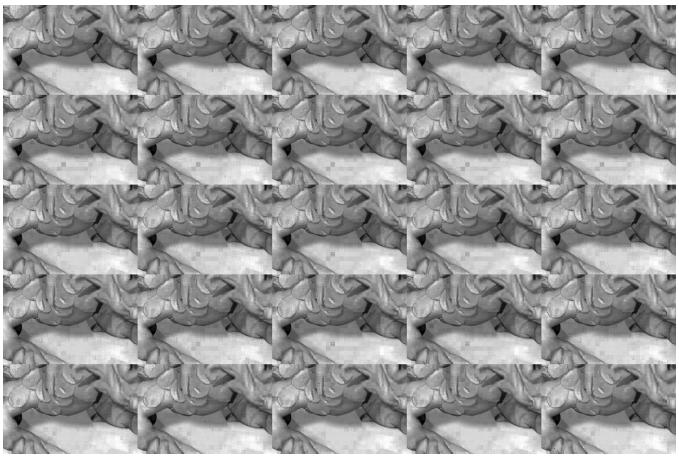

(b) Tensor NLLR

Fig. 10. Reconstructed light fields

are sufficiently close in ℓ_2 distance (unlike the distribution shown in Figure 1) for the ℓ_2-NLLR model to be accurate. Fast LRSS worked with complex images, while clustering in ℓ_2-NLLR used magnitude images to remove phase variation, and therefore Fast LRSS may be beneficial for MRI applications where useful information is encoded in phase.

B. Tensor Reconstruction Comparison Results

Although the matrix NLLR reconstructions resulted in higher RMS errors, many of the artifacts could probably be ameliorated with overlapping blocks or by ensuring that *all* similar blocks within a certain distance appear in the same cluster, instead of just 6. Also, as mentioned previously, patches should be ensured to appear in multiple clusters to increase the robustness of the reconstruction.

In our results, we showed that tensor reconstruction can achieve good results without specific assumptions of redundancy in certain dimensions. In the future, it would be useful to figure out which dimensions contribute the most to light field reconstruction using low tensor rank, which could more clearly identify the benefits of tensor vs. matrix NLLR reconstruction. Other future work includes optimizing the MATLAB code to run faster in order to increase the number of neighbors in clusters. More patches in clusters and more overlap between the clusters could be beneficial to increasing robustness of the algorithm. Subspace clustering as described in the first section of the report could also be applied to tensor reconstruction. In addition, currently, the proposed algorithm requires some tuning in regards to the starting estimate rank. It would be useful to include a method to estimate the starting rank estimate of the desired low rank tensor cluster. Finally, it would be ideal to develop a convex formulation of tensor completion, possibly involving other decompositions such as variations of Tucker decomposition, while avoiding tensor unfolding.

V. Summary and Conclusion

Many image reconstruction and analysis tasks make use of linear dependencies, which can be formulated as low rank matrix or tensor approximations. Images have self-similar structure, and linear dependencies may be non-local. In this project, we have shown through denoising and MRI applications that applying subspace clustering to image patches is a natural way to learn non-local low rank structure. In addition, for multidimensional imaging applications such as Light field Imaging, extending the non-local low rank framework to enforcing tensor low rank is a promising method to deal with problem-dependent data reformatting.

Relation to Previous Research

MRI acquisitions were not specifically done for the purpose of this project and can be considered other research. Specifically for the project, NLLR reconstruction was formulated and implemented in Matlab using functions for SVT, ADMM, SURE, and block-wise Matricization that were written for other purposes and required minimal modifications. Results were mostly anecdotal for the purpose of investigating ideas that would possibly be followed up with research. Putting code together, extending it from local to non-local low rank, writing all clustering functions, and non-local analysis comprised most of the effort in software development. This report was written for the class project is not a recycled conference paper. Open questions about NLLR were discussed with Felix and Prof. Wetzstein, but all novel contributions related to tensors and subspace clustering were original.

Team project split

Evan

1) Literature Review: Non-local low rank regularization approaches, subspace clustering methods, hashing of patches, gleaning wisdom from Felix Heide.
2) Theory/Algorithms: subspace clustering, analysis and formulation of non-local low rank, connection to subspace clustering.
3) Software: matricization, prox functions for matrix problems, minor modifications of an existing ADMM function, implementation of SSC and LRSS, compiling existing algorithms such as spectral clustering and nearest neighbor search.

4) Applications/Experiments: CS-MRI and image denoising.

Tiffany

1) Literature Review: Tensor Decomposition, Tensor Completion, Lightfield Imaging Completion, Lightfield Inpainting, review of applications in computational imaging (HSI, computational illumination, computational photography, light field imaging)
2) Theory/Algorithms: tensor decompositions for images, tensors thresholding
3) Software: prox functions for tensor problems, clustering and grouping operator for tensors
4) Applications/Experiments: Light field image reconstruction

ACKNOWLEDGMENT

The authors acknowledge Felix Heide and Gordon Wetzstein for helpful discussions.

REFERENCES

[1] J. Wright, Y. Ma, J. Mairal, G. Sapiro, T. S. Huang, and S. Yan, "Sparse representation for computer vision and pattern recognition," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1031–1044, 2010.

[2] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-d transform-domain collaborative filtering," *Image Processing, IEEE Transactions on*, vol. 16, no. 8, pp. 2080–2095, 2007.

[3] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 2. IEEE, 2005, pp. 60–65.

[4] W. Dong, G. Shi, X. Li, Y. Ma, and F. Huang, "Compressive sensing via nonlocal low-rank regularization," *Image Processing, IEEE Transactions on*, vol. 23, no. 8, pp. 3618–3632, Aug 2014.

[5] J. D. Trzasko, "Exploiting local low-rank structure in higher-dimensional mri applications," in *SPIE Optical Engineering+ Applications*. International Society for Optics and Photonics, 2013, pp. 885 821–885 821.

[6] J. Trzasko, A. Manduca, and E. Borisch, "Local versus global low-rank promotion in dynamic mri series reconstruction," in *Proc. Int. Symp. Magn. Reson. Med*, 2011, p. 4371.

[7] R. Otazo, E. Candès, and D. K. Sodickson, "Low-rank plus sparse matrix decomposition for accelerated dynamic mri with separation of background and dynamic components," *Magnetic Resonance in Medicine*, 2014.

[8] T. Zhang, J. Y. Cheng, A. G. Potnick, R. A. Barth, M. T. Alley, M. Uecker, M. Lustig, J. M. Pauly, and S. S. Vasanawala, "Fast pediatric 3d free-breathing abdominal dynamic contrast enhanced mri with high spatiotemporal resolution," *Journal of Magnetic Resonance Imaging*, 2013.

[9] R. Vidal and R. Hartley, "Motion segmentation with missing data using powerfactorization and gpca," in *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, vol. 2, June 2004, pp. II–310–II–316 Vol.2.

[10] R. Vidal, "A tutorial on subspace clustering," *IEEE Signal Processing Magazine*, vol. 28, no. 2, pp. 52–68, 2010.

[11] E. Elhamifar and R. Vidal, "Sparse subspace clustering," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 2009, pp. 2790–2797.

[12] M. Soltanolkotabi, E. Elhamifar, E. J. Candes *et al.*, "Robust subspace clustering," *The Annals of Statistics*, vol. 42, no. 2, pp. 669–699, 2014.

[13] A. Y. Ng, M. I. Jordan, Y. Weiss *et al.*, "On spectral clustering: Analysis and an algorithm," *Advances in neural information processing systems*, vol. 2, pp. 849–856, 2002.

[14] X. Zhang, F. Sun, G. Liu, and Y. Ma, "Fast low-rank subspace segmentation," *Knowledge and Data Engineering, IEEE Transactions on*, vol. 26, no. 5, pp. 1293–1297, 2014.

[15] D. Goldfarb and Z. Qin, "Robust low-rank tensor recovery: Models and algorithms," *SIAM J. Matrix Anal. Appl*, 2013.

[16] Y. Liu, F. Shang, H. Cheng, J. Cheng, and H. Tong, "Factor matrix trace norm minimization for low-rank tensor completion," in *Procc of the 2014 SIAM International Conference on Data Mining*, 2014.

[17] L. Yatziv, G. Sapiro, and M. Levoy, "Lightfield completion," in *Proc. IEEE Int. Conf. Image Processing*, vol. 3, 2004, pp. 1787–1790.

[18] W. Li, L. Zhao, Z. Lin, D. Xu, and D. Lu, "Non-local image inpainting using low-rank matrix completion," *Computer Graphics Forum*, 2014.

[19] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, no. 3, pp. 455–500, 2009.

[20] J. Liu, P. Musialski, P. Wonka, and J. Ye, "Tensor completion for estimating missing values in visual data," in *ICCV*, 2009, pp. 2114–2121.

[21] B. W. Bader and T. G. Kolda, "Tensor decompositions, the matlab tensor toolbox, and applications to data analysis [powerpoint slides]," *Retrieved from https://www.ima.umn.edu/industrial/2006-2007/kolda/kolda.pdf*, 2007.

[22] 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*, vol. 3, no. 1, pp. 1–122, 2011.

[23] E. J. Candes, C. A. Sing-Long, and J. D. Trzasko, "Unbiased risk estimates for singular value thresholding and spectral estimators," *Signal Processing, IEEE Transactions on*, vol. 61, no. 19, pp. 4643–4657, 2013.

[24] E. Levine, B. Daniel, B. Hargreaves, and S. MS, "Validation of reduced view-sharing compressed sensing reconstruction for dce-mri with variable flip angle acquisition," in *Proceedings of the 23rd Annual Meeting of ISMRM, Toronto, Canada*, 2015.

[25] E. Levine, M. Saranathan, and H. BA, "Complementary poisson-disc sampling," in *Proceedings of the 22nd Annual Meeting of ISMRM, Milan, Italy*, 2014.

[26] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and M. Lustig, "Espiritan eigenvalue approach to autocalibrating parallel mri: where sense meets grappa," *Magnetic Resonance in Medicine*, vol. 71, no. 3, pp. 990–1001, 2014.