



Published in final edited form as:

Proc Int Conf Image Proc. 2016 September ; 2016: 963–967. doi:10.1109/icip.2016.7532500.

STRUCTURED LOW-RANK RECOVERY OF PIECEWISE CONSTANT SIGNALS WITH PERFORMANCE GUARANTEES

Greg Ongie^{*}, Sampurna Biswas[†], Mathews Jacob[†]

^{*} Department of Mathematics, University of Iowa, IA, USA.

[†] Department of Electrical and Computer Engineering, University of Iowa, IA, USA

Abstract

We derive theoretical guarantees for the exact recovery of piecewise constant two-dimensional images from a minimal number of non-uniform Fourier samples using a convex matrix completion algorithm. We assume the discontinuities of the image are localized to the zero level-set of a bandlimited function, which induces certain linear dependencies in Fourier domain, such that a multifold Toeplitz matrix built from the Fourier data is known to be low-rank. The recovery algorithm arranges the known Fourier samples into the structured matrix then attempts recovery of the missing Fourier data by minimizing the nuclear norm subject to structure and data constraints. This work adapts results by Chen and Chi on the recovery of isolated Diracs via nuclear norm minimization of a similar multifold Hankel structure. We show that exact recovery is possible with high probability when the bandlimited function describing the edge set satisfies an incoherency property. Finally, we demonstrate the algorithm on the recovery of undersampled MRI data.

Keywords

Structured Low-Rank Matrix Completion; Annihilating Filter Method; Finite Rate of Innovation; Compressed Sensing; MRI

1. INTRODUCTION

The recovery of a linear combination of exponentials from their few uniform samples is a classical problem in signal processing with extensive applications. Prony's method, or one of its robust variants, attempts to recover the signal by estimating an annihilating polynomial whose zeros correspond to the frequency of the exponentials. The finite rate of innovation (FRI) framework [1] extended these methods to recover more general signals that reduce to a sparse linear combination of Dirac delta functions under an appropriate transformation (e.g., differential operators, convolution). Recently, several authors have further extended FRI methods to recover such signals from their non-uniform Fourier samples [2–6] by exploiting the low-rank structure of an enhanced matrix (e.g., Hankel matrix in 1-D). Performance guarantees do exist when the transform is an identity and when the Diracs are well-separated [2].

The above signal models have limited flexibility in exploiting the extensive additional structure present in many multidimensional imaging problems. Specifically, the edges in multidimensional images are connected and can be modeled as smooth curves or surfaces.

We have recently introduced a novel framework to recover piecewise polynomial images, whose edges are localized to smooth curves, from their uniform [7,8] and non-uniform [6] Fourier samples; this work generalizes a recent extension of the FRI framework to curves [9]. We model the piecewise smooth signal as having partial derivatives that vanish outside the zero level-set of a bandlimited function. This relation translates to an annihilation condition involving the uniform Fourier samples of the partial derivatives, which can be compactly represented as the multiplication of a specific structured matrix with the Fourier coefficients of the bandlimited function. Our earlier work has shown that the structured matrix is low-rank, and we used this property to recover the signal from its non-uniform Fourier samples with good performance. Efficient algorithms that work on the original signal samples rather than the structured high-dimensional matrix also were introduced [10]. We observe the signal models in [2,3,5] do not include the class of signals considered in this work.

The main focus of this work is to introduce theoretical guarantees on the recovery of piecewise constant signals, whose discontinuities are localized to zero level-sets of bandlimited functions, from non-uniform Fourier samples. Since such signals cannot be expressed as a finite linear combination of isolated Diracs, the recovery guarantees in [2] cannot be directly extended to our setting. Specifically, the theory in [2] relies heavily on an explicit factorization of the enhanced matrix (e.g., Vandermonde factorization of a Hankel matrix in the 1-D case), which is only available when the number of discontinuities are finite and well separated. Instead, we give a new description of the row and column subspace of the structured matrix, which allow us to derive incoherence measures based solely on properties of the bandlimited function describing the edge set of the image.

2. THEORY

2.1. Signal Model: 2-D Piecewise Constant Images

We consider the recovery of a *piecewise constant* function

$$f(\mathbf{r}) = \sum_{i=1}^N a_i \chi_{\Omega_i}(\mathbf{r}), \quad \text{for all } \mathbf{r} = (x, y) \in [0, 1]^2, \quad (1)$$

where $a_i \in \mathbb{C}$, and χ_{Ω} denotes the characteristic function of the set Ω . We assume the Fourier samples $\hat{f}[\mathbf{k}]$ specified by

$$\hat{f}[\mathbf{k}] = \int_{[0,1]^2} f(\mathbf{r}) e^{-j2\pi\mathbf{k} \cdot \mathbf{r}}, \quad \mathbf{k} \in \mathbb{Z}^2, \quad (2)$$

are available at a subset of non-uniform locations $\mathbf{k} \in \Theta$ belonging to a rectangular set of uniform sampling locations $\Gamma \subset \mathbb{Z}^2$ in Fourier domain.

We further assume that the edge-set of the image, specified by $E := \cup_i \Omega_i$, to be the zero-set of a 2-D bandlimited trigonometric polynomial:

$$\mu_0(\mathbf{r}) = \sum_{\mathbf{k} \in \Lambda_0} c[\mathbf{k}]e^{j2\pi\mathbf{k} \cdot \mathbf{r}}, \quad \forall \mathbf{r} \in [0, 1]^2, \tag{3}$$

where the coefficients $c[\mathbf{k}] \in \mathbb{C}$, and Λ_0 is a rectangular subset of Γ . Here we assume μ_0 is the unique minimal degree trigonometric polynomial such that $E = \{\mu_0 = 0\}$, where the degree is defined by the rectangular dimensions of the Fourier support Λ_0 . We have shown in [8] that when $\mu(\mathbf{r})$ is any bandlimited function that vanishes on the edgeset E , the gradient $\nabla f = (x f, y f)$ satisfies the property

$$\mu \nabla f = 0, \tag{4}$$

in the distributional sense. See Fig. 1 for an illustration when the number of regions $N=1$. Note that among all functions bandlimited to Λ_0 in Fourier domain, $\mu = \mu_0$ is the only one in this class that satisfies (4). However, if we consider μ that are bandlimited to a larger rectangular set Λ_1 with $\Lambda_0 \subset \Lambda_1 \subset \Gamma$, then we have shown that all μ satisfying (4) are a multiple of μ_0 [8]. The spatial domain relation (4) translates directly to the following vector annihilation relation in the Fourier domain:

$$\sum_{\mathbf{k} \in \Lambda_1} \widehat{\nabla f}[\ell - \mathbf{k}] \hat{\mu}[\mathbf{k}] = \mathbf{0}, \quad \forall \ell \in \Lambda_2. \tag{5}$$

Here $\widehat{\nabla f}[\mathbf{k}] = j2\pi(k_1 \hat{f}[\mathbf{k}], k_2 \hat{f}[\mathbf{k}])$ for $\mathbf{k} = (k_1, k_2)$, and $\Lambda_2 \subset \mathbb{Z}^2$ is any rectangular set on which the convolutions between $\widehat{\nabla f}$ and $\hat{\mu}$ is valid. Note that the 2-D convolution between two filters support limited to rectangular sets Λ_1 and Λ_2 is support limited to the dilation of Λ_1 by Λ_2 , which we denote by $\Lambda_1 * \Lambda_2$. Since $\hat{\mu}$ is bandlimited to Λ_1 , when using samples of $\widehat{\nabla f}$ within Γ , we require $\Lambda_1 * \Lambda_2 = \Gamma$; see Fig. 2.

The Fourier domain annihilation relations (5) can be compactly represented in matrix form as

$$\mathcal{T}(\hat{f})\mathbf{h} = \begin{bmatrix} \mathcal{T}_1(\hat{f}) \\ \mathcal{T}_2(\hat{f}) \end{bmatrix} \mathbf{h} = \mathbf{0}, \tag{6}$$

where $\mathcal{T}_i(\hat{f}) \in \mathbb{C}^{|\Lambda_2| \times |\Lambda_1|}$, $i = 1, 2$, are matrices corresponding to the discrete 2-D convolution of $k_i \hat{f}[\mathbf{k}]$, $\mathbf{k} \in \Gamma$, (omitting the irrelevant factor $j2\pi$) with a filter supported on Λ_1 , with output restricted to the index set Λ_2 . Here we use \mathbf{h} to denote the vectorized version of a filter $h[\mathbf{k}]$, $\mathbf{k} \in \Lambda_1$. By our previous observations, the solutions \mathbf{h} to (6) are given by the Fourier coefficients of a multiple of the minimal polynomial. Hence if the filter support Λ_1 is larger than the minimal filter support Λ_0 , $\mathcal{T}(\hat{f})$ has a large nullspace and is low-rank. Specifically, in [8] we proved the following:

Proposition 1.—*Suppose $\mathcal{T}(\hat{f})$ is built with filter size $\Lambda_1 \supseteq \Lambda_0$ satisfying $\Gamma \supseteq 2\Lambda_1 * \Lambda_0$, then*

$$R := \text{rank } \mathcal{T}(\hat{f}) = |\Lambda_1| - |\Lambda_1 : \Lambda_0|$$

where $|\Lambda_1|$ is the number of indices in Λ_1 and $|\Lambda_1 : \Lambda_0|$ is the number of integer shifts of Λ_0 contained in Λ_1 .

2.2. Recovery from non-uniform Fourier samples

Since the matrix $\mathcal{T}(\hat{f})$ is low-rank, we propose to recover the signal from its noiseless Fourier samples as the convex optimization problem:

$$\text{minimize}_{\hat{g}} \|\mathcal{T}(\hat{g})\|_* \text{ subject to } \hat{g}[\mathbf{k}] = \hat{f}[\mathbf{k}]; \mathbf{k} \in \Theta \quad (7)$$

where $\|\cdot\|_*$ denotes the nuclear norm, i.e., the sum of the singular values. To aid in our analysis, we will now reformulate the recovery of \hat{f} as a matrix recovery problem using projection operators in the lifted matrix domain. We define basis matrices

$$\mathbf{A}_{\mathbf{k}} = \begin{bmatrix} \mathbf{A}_{1, \mathbf{k}} \\ \mathbf{A}_{2, \mathbf{k}} \end{bmatrix} \in \mathbb{C}^{2|\Lambda_2| \times |\Lambda_1|}, \text{ for all } \mathbf{k} \in \Gamma, \text{ where}$$

$$(\mathbf{A}_{i, \mathbf{k}})_{\alpha, \beta} = \begin{cases} \frac{k_i}{|\mathbf{k}| \sqrt{|\omega_i(\mathbf{k})|}}, & \text{if } (\alpha, \beta) \in \omega_i(\mathbf{k}) \\ 0 & \text{else} \end{cases} \quad (8)$$

for $i = 1, 2$. Here $\omega_i(\mathbf{k})$ is the set of locations of the matrix $\mathcal{T}(\hat{f})$ containing copies of $k_i \widehat{f}[\mathbf{k}]$. Note that the set $\{\mathbf{A}_{\mathbf{k}}\}_{\mathbf{k} \in \Gamma}$ forms an orthonormal basis for the space of structured matrices defined by the lifting \mathcal{T} . Specifically, for any set of coefficients $\{\hat{g}[\mathbf{k}]\}_{\mathbf{k} \in \Gamma}$ we can expand the matrix $\mathbf{X} = \mathcal{T}(\hat{g})$ as $\mathbf{X} = \sum_{\mathbf{k}} \langle \mathbf{X}, \mathbf{A}_{\mathbf{k}} \rangle \mathbf{A}_{\mathbf{k}}$, where $\langle \mathbf{X}, \mathbf{A}_{\mathbf{k}} \rangle = \hat{g}[\mathbf{k}]$. Using these basis matrices, we define the following operators in the lifted domain:

$$\mathcal{A}_{\Theta}(\mathbf{X}) = \sum_{\mathbf{k} \in \Theta} \langle \mathbf{A}_{\mathbf{k}}, \mathbf{X} \rangle \mathbf{A}_{\mathbf{k}} \quad (9)$$

$$\mathcal{A}^{\perp}(\mathbf{X}) = \mathcal{T} - \sum_{\mathbf{k} \in \Gamma} \langle \mathbf{A}_{\mathbf{k}}, \mathbf{X} \rangle \mathbf{A}_{\mathbf{k}} \quad (10)$$

$$\mathcal{Q}_{\Theta} = \frac{|\Gamma|}{|\Theta|} \mathcal{A}_{\Theta} + \mathcal{A}^{\perp} \quad (11)$$

The constants in (11) are chosen so that $\mathbb{E}[\mathcal{Q}_{\Theta}] = \mathcal{T}$. Using these definitions, we rewrite (7) as

$$\text{minimize}_{\mathbf{X}} \|\mathbf{X}\|_* \text{ subject to } \mathcal{Q}_{\Theta}(\mathbf{X}) = \mathcal{Q}_{\Theta}(\mathcal{T}(\hat{f})) \quad (12)$$

Several authors have shown that the performance of low-rank matrix recovery by nuclear norm minimization is dependent on the *incoherence* of the sampling basis with respect to the matrix to be recovered [2, 11]. Towards this end, we introduce two incoherence measures associated with $\mathcal{T}(\hat{f})$ defined solely in terms of the edge-set polynomial μ_0 . In the following, we set D_{Λ_1} to be the 2-D Dirichlet kernel supported on Λ_1 , i.e., the function such that $\widehat{D_{\Lambda_1}}[\mathbf{k}] = 1$ for all $\mathbf{k} \in \Lambda_1$ and zero otherwise. For any collection of N points $P = \{\mathbf{r}_1, \dots, \mathbf{r}_N\} \subset [0, 1]^2$, we define the $N \times N$ Gram matrix $\mathbf{G}(P)$ by $(\mathbf{G}(P))_{i,j} = D_{\Lambda_1}(\mathbf{r}_i - \mathbf{r}_j)$.

Definition 2.—Let μ_0 be a trigonometric polynomial bandlimited to Λ_0 (see (3)), and set $R = |\Lambda_1| - |\Lambda_1 \cap \Lambda_0|$. Define the incoherence measure ρ_1 by

$$\frac{1}{\rho_1} = \max_{\substack{P \subset \{\mu_0 = 0\} \\ |P| = R}} \sigma_{\min}[\mathbf{G}(P)] \tag{13}$$

where $\sigma_{\min}[\mathbf{G}(P)]$ is the minimum singular value of $\mathbf{G}(P)$.

Put in words, among all possible arrangements of R points along the edge-set $\{\mu_0 = 0\}$, we seek the arrangement that gives the best conditioning of the matrix $\mathbf{G}(P)$, and call the resulting condition number ρ_1 . Intuitively, the optimal arrangement will maximize the minimum separation distance among the R points, and ρ_1 can be thought of as a measure of this geometric property. In particular, having any edges that enclose a small area will give a high ρ_1 .

Additionally, our results rely on another incoherence measure related to properties of the gradient of μ_0 :

Definition 3.—Let μ_0 be a trigonometric polynomial bandlimited to Λ_0 (see (3)). Normalize μ_0 by $\int_{[0, 1]^2} |\nabla \mu_0(\mathbf{r})|^2 d\mathbf{r} = 1$. Define the incoherence measure ρ_2 by

$$\frac{1}{\rho_2} = \frac{\min_{\gamma \in B_{\Lambda_1}} \int_{[0, 1]^2} |\gamma(\mathbf{r})|^2 |\nabla \mu_0(\mathbf{r})|^2 d\mathbf{r}}{\|\gamma\|_2 = 1} \frac{1}{\|\widehat{\nabla \mu_0}\|_1^2} \tag{14}$$

where B_{Λ_1} denotes the space of all trigonometric polynomials bandlimited to Λ_1 , and $\|\widehat{\nabla \mu_0}\|_1 = \sum_{\mathbf{k} \in \Lambda_0} |\widehat{\nabla \mu_0}[\mathbf{k}]|$.

Note that ρ_2 will be large when $|\nabla \mu_0|$ has several zeros, or equivalently, when μ_0 has several critical points. Since μ_0 must contain a critical point in every region defined by the complement of the edge-set, ρ_2 will be large when the image has several distinct regions.

Now we state our main result:

Theorem 4.—Let f be specified by (1), whose edge-set is described by the zero-set of the trigonometric polynomial μ_0 bandlimited to Λ_0 (see (3)) with associated incoherence

measures ρ_1 and ρ_2 . Let $\Theta \subset \Gamma$ be an index set drawn uniformly at random within Γ . Then there exists a universal constant $c > 0$ such that the solution to (12) is exact ($\mathbf{X} = T(\hat{f})$) with probability exceeding $1 - |\Gamma|^{-2}$, provided

$$|\Theta| > c\rho_1\rho_2Rc_s\log^4|\Gamma|. \tag{15}$$

where $R = \text{rank } \mathcal{T}(\hat{f})$ and $c_s = |\Gamma|/|\Lambda_1|$.

Following the approach in [2], we can prove this result by constructing an approximate dual certificate using the well-known golfing scheme of [11]. The adaptation of the proof in [2] to the measurement operator (11) is straightforward, and these details are omitted for brevity. The essential difference of the above result and [2] is the characterization of the incoherency measures. The approach in [2] relies on an explicit low-rank factorization of the lifted matrix in terms of Vandermonde-like matrices, which is not available in our setting. Instead, we characterize the row and column spaces of the matrix and use it to prove the above result.

2.3. Row and column spaces of $\mathcal{T}(\hat{f})$ and incoherence

Define \mathcal{P}_U and \mathcal{P}_V to be the orthogonal projections onto the column space and row space of $\mathcal{T}(\hat{f})$, respectively, i.e., if $\mathcal{T}(\hat{f}) = \mathbf{U}\Sigma\mathbf{V}^*$ is the rank- R singular value decomposition then $\mathcal{P}_U\mathbf{X} = \mathbf{U}\mathbf{U}^*\mathbf{X}$, $\mathcal{P}_V\mathbf{X} = \mathbf{X}\mathbf{V}\mathbf{V}^*$. One can show it is possible to construct an approximate dual certificate with high probability [2, 11], provided we can uniformly bound the norms of the projections $\mathcal{P}_U\mathbf{A}_k$ and $\mathcal{P}_V\mathbf{A}_k$. The following proposition shows these norms can be controlled by the incoherence measures introduced in (13) and (14):

Proposition 5.—Consider $\mathcal{T}(\hat{f})$ of rank R corresponding to a piecewise constant function f whose edge set coincides with the zero set of μ_0 , let ρ_1 and ρ_2 be the incoherency measures of μ_0 , and set $c_s = |\Gamma|/|\Lambda_1|$. Then we have

$$\max_{\mathbf{k} \in \Gamma} \{ \|\mathcal{P}_U\mathbf{A}_k\|_F^2, \|\mathcal{P}_V\mathbf{A}_k\|_F^2 \} \leq \frac{\rho_1\rho_2Rc_s}{|\Gamma|}. \tag{16}$$

The proof relies on the following basis representations for the row and column spaces of $\mathcal{T}(\hat{f})$:

Lemma 6.—Choose any $N = |\Lambda_1| + |\Lambda_0| - |\Lambda_1 \cap \Lambda_0|$ points $\{\mathbf{r}_1, \dots, \mathbf{r}_N\} \subset \{\mu_0 = 0\}$, and define B to be the collection of filters $\{\mathbf{h}_1, \dots, \mathbf{h}_N\}$ where each $\mathbf{h}_i \in \mathbb{C}^{|\Lambda|}$ are the Fourier coefficients of the translated Dirichlet kernel $D_{\Lambda_1}(\mathbf{r} - \mathbf{r}_i)$. Then there exists a subset of $R = \text{rank } \mathcal{T}(\hat{f})$ elements from B that is a basis for the row space of $\mathcal{T}(\hat{f})$.

Lemma 7.—Let $B_R = \{\mathbf{d}_i\}_{i=1}^R$ be any basis of the row space of $\mathcal{T}(\hat{f})$, and set $\widehat{\nabla}\mu_0 = (\mathbf{c}_x, \mathbf{c}_y) \in \mathbb{C}^{2|\Lambda_0|}$. Then $B_L = \{(\mathbf{d}_i * \mathbf{c}_x, \mathbf{d}_i * \mathbf{c}_y)\}_{i=1}^R$ is a basis of the column space of $\mathcal{T}(\hat{f})$.

3. EXPERIMENTS

In Fig. 3 we demonstrate the recovery of a synthetic piecewise constant phantom [12] (181×181 sampling grid, single channel) from 20-fold variable density random Fourier samples using the structured low-rank matrix completion approach (7). We solve (7) using a singular value thresholding approach proposed in [6]. The filter size was set to 33×25 . Compared with total variation (TV) minimization, the proposed structured low-rank approach more accurately recovers the original piecewise constant regions.

Additionally, in Fig. 4 we demonstrate the structured low-rank approach on the recovery real MR data (255×255 sampling grid with 4 receiver coils, coil-compressed to a single channel) from 2-fold uniform random undersampling, using a filter size of 45×45 . Due to the problem size, the formulation (7) is difficult to solve via singular value thresholding. Instead, we make use of the recently proposed GIRAF algorithm [10] which solves an approximated version of (7). The result shows similar benefit over a TV regularized recovery in its ability to preserve fine details and strong edges.

4. DISCUSSION AND CONCLUSION

We derived a performance guarantee for the recovery of piecewise constant images from non-uniform Fourier samples by a structured matrix completion. This was achieved by adapting results in [2] to the case of a low-rank multifold Toeplitz structure with an additional weighting scheme. We also define new incoherence measures that rely only on properties of the minimal annihilating polynomial whose zero-set encodes the edges of the image.

While in the present work we only consider noiseless ideal samples, in future work we intend to derive guarantees for robust recovery in the presence of noise and model-mismatch. Additionally, it would be interesting to adapt our results to a wider variety of sampling distributions, and to identify the optimal sampling strategy for signals belonging to our image model.

Acknowledgments

This work is supported by grants NIH 1R01EB019961-01A1, NSF CCF-1116067, ONR N00014-13-1-0202, and ACS RSG-11-267-01-CCE.

5. REFERENCES

- [1]. Vetterli M, Marziliano P, and Blu T, "Sampling signals with finite rate of innovation," *Signal Processing, IEEE Transactions on*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [2]. Chen Y and Chi Y, "Robust spectral compressed sensing via structured matrix completion," *Information Theory, IEEE Transactions on*, vol. 60, no. 10, pp. 6576–6601, 2014.
- [3]. Jin KH, Lee D, and Ye JC, "A general framework for compressed sensing and parallel MRI using annihilating filter based low-rank hankel matrix," *arXiv preprint arXiv:1504.00532*, 2015.
- [4]. Jin KH, Lee D, and Ye JC, "A novel k-space annihilating filter method for unification between compressed sensing and parallel mri," in *International Symposium on Biomedical Imaging (ISBI)*, 4 2015.

- [5]. Haldar JP, “Low-rank modeling of local k-space neighborhoods (LORAKS) for constrained MRI.” *Medical Imaging, IEEE Transactions on*, vol. 33, no. 3, pp. 668–681, 2014.
- [6]. Ongie G and Jacob M, “Recovery of piecewise smooth images from few fourier samples,” in *Sampling Theory and Applications (SampTA)*, 5 2015, pp. 543–547.
- [7]. Ongie G and Jacob M, “Super-resolution MRI using finite rate of innovation curves,” *IEEE International Symposium on Biomedical Imaging (ISBI)*, 4 2015.
- [8]. Ongie G and Jacob M, “Off-the-grid recovery of piecewise constant images from few fourier samples,” in *arXiv:1510.00384*, 5 2015, pp. 543–547.
- [9]. Pan H, Blu T, and Dragotti PL, “Sampling curves with finite rate of innovation,” *Signal Processing, IEEE Transactions on*, vol. 62, no. 2, 2014.
- [10]. Ongie G and Jacob M, “A fast algorithm for structured low-rank matrix recovery with applications to undersampled MRI reconstruction,” *IEEE International Symposium on Biomedical Imaging (ISBI)*, 4 2016.
- [11]. Gross D, “Recovering low-rank matrices from few coefficients in any basis,” *Information Theory, IEEE Transactions on*, vol. 57, no. 3, pp. 1548–1566, 2011.
- [12]. Guerquin-Kern M, Lejeune L, Pruessmann KP, and Unser M, “Realistic analytical phantoms for parallel magnetic resonance imaging,” *Medical Imaging, IEEE Transactions on*, vol. 31, no. 3, pp. 626–636, 2012.

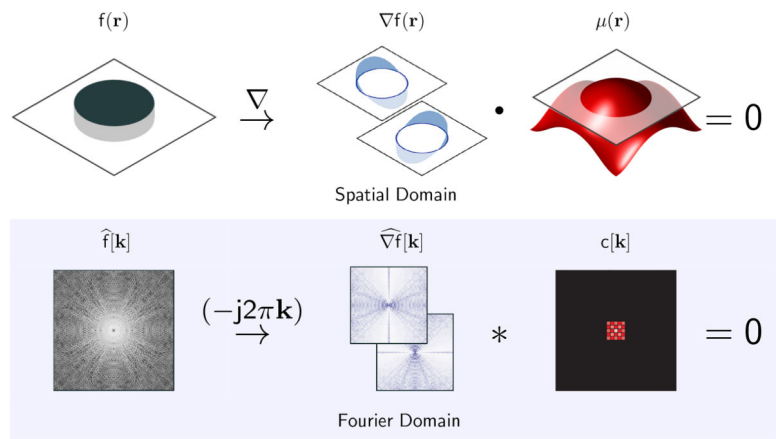


Fig. 1. Annihilation of a piecewise constant function in the spatial (top) and Fourier (bottom) domain.

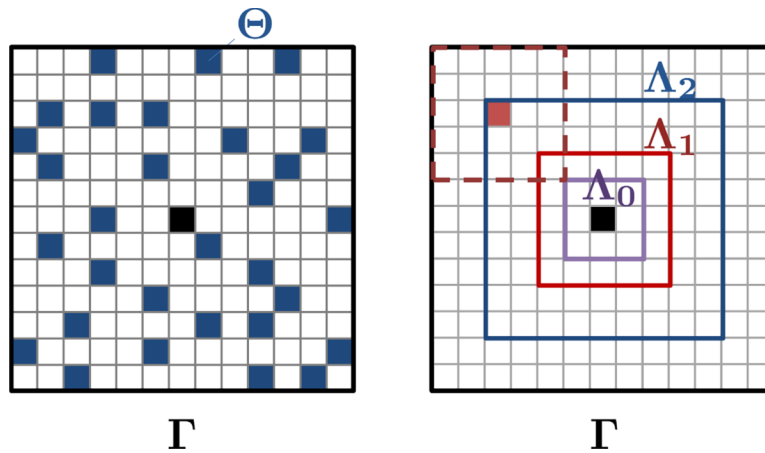


Fig. 2. Fourier domain support sets used in constructing the structured matrix $\mathcal{T}(\hat{f})$. The grid Γ represent a rectangular sampling window in \mathbb{Z}^2 (index $(0, 0)$ marked in black). (left) Θ is a collection of uniform random sampling locations. (right) Λ_0 is the Fourier support of the minimal annihilating polynomial, Λ_1 is the assumed filter size used in constructing $\mathcal{T}(\hat{f})$, and Λ_2 is the set of valid convolution locations satisfying $\Lambda_1 * \Lambda_2 = \Gamma$. Note that $\mathcal{T}(\hat{f})$ has dimensions $2|\Lambda_2| \times |\Lambda_1|$.

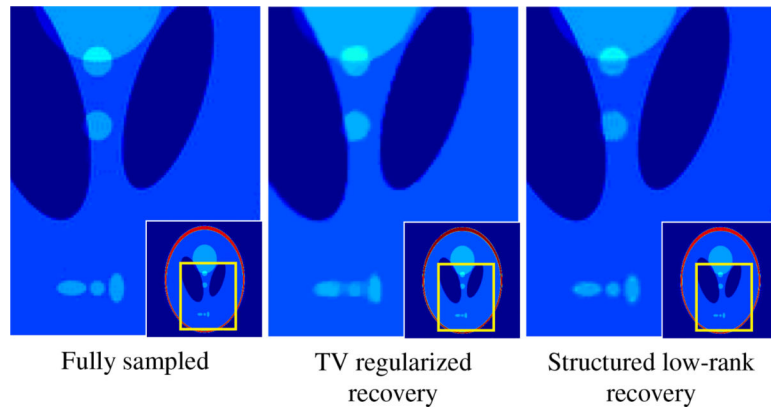


Fig. 3. Recovery of synthetic MRI data from 20-fold variable density undersampling.

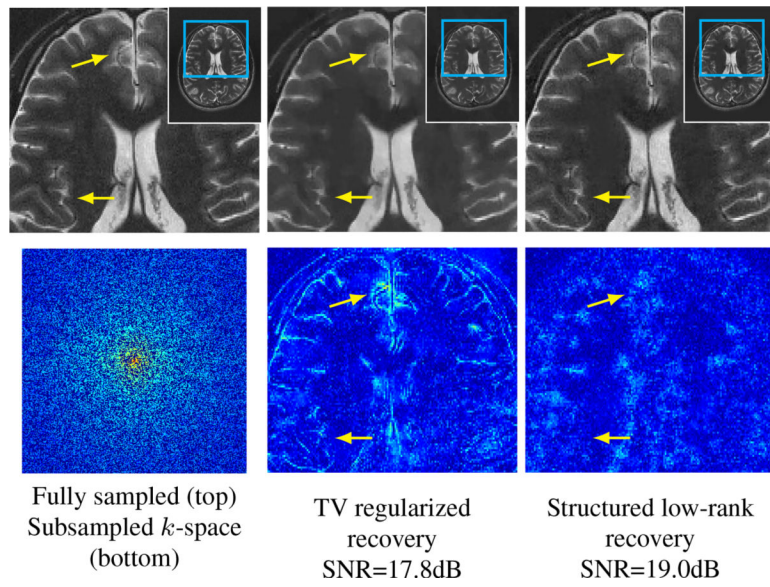


Fig. 4. Recovery of real MRI data from 2-fold random uniform undersampling. Error images shown below.