



HHS Public Access

Author manuscript

Proc IEEE Int Symp Biomed Imaging. Author manuscript; available in PMC 2021 February 21.

Published in final edited form as:

Proc IEEE Int Symp Biomed Imaging. 2018 April ; 2018: 1272–1275. doi:10.1109/isbi.2018.8363803.

RECOVERY OF POINT CLOUDS ON SURFACES: APPLICATION TO IMAGE RECONSTRUCTION

Sunrita Poddar, Mathews Jacob

Department of Electrical and Computer Engineering, University of Iowa, IA, USA

Abstract

We introduce a framework for the recovery of points on a smooth surface in high-dimensional space, with application to dynamic imaging. We assume the surface to be the zero-level set of a bandlimited function. We show that the exponential maps of the points on the surface satisfy annihilation relations, implying that they lie in a finite dimensional subspace. We rely on nuclear norm minimization of the maps to recover the points from noisy and undersampled measurements. Since this direct approach suffers from the curse of dimensionality, we introduce an iterative reweighted algorithm that uses the "kernel trick". The resulting algorithm has similarities to iterative algorithms used in graph signal processing (GSP); this framework can be seen as a continuous domain alternative to discrete GSP theory. The use of the algorithm in recovering free breathing and ungated cardiac data shows the potential of this framework in practical applications.

Keywords

machine learning; kernels; superresolution; denoising; dynamic MRI

1. INTRODUCTION

The recovery of signals that lie on a manifold/surface has received extensive attention in the recent years. These methods model the high-dimensional data as points localized to low-dimensional manifolds. For example, patch-based image processing methods such as BM3D implicitly use the structure of the patch manifold [1, 2, 3], while we [4, 5] and others [6, 7] have recently used the manifold structure of images in a dynamic time series. Manifold embedding methods are also widely used in machine learning as visualization tools.

The main focus of this paper is to introduce a continuous domain perspective for denoising/regularization of points (e.g patches or images), which are assumed to be drawn from a smooth surface in a very high dimensional space. This surface can be represented as the zero-level set of a bandlimited potential function. We note that the level-set model can account for a large variety of surfaces and is widely used in image processing. We observe that the potential function is zero for any point on the surface, which we term as an annihilation relation. We show that the annihilation relation can be expressed as a weighted linear combination of exponential feature maps of the point. The dimension of the feature map is equal to the bandwidth of the potential function. When the bandwidth is overestimated, there are multiple such annihilation relations, suggesting that the exponential feature maps of the points on the surface lie in a finite dimensional space. We show that the

finite dimensional nature of the maps translates to a low-rank kernel matrix, computed from the points using a shift invariant kernel function such as the Dirichlet and Gaussian kernels. This demonstrates the direct link between superresolution signal recovery [8] and kernel methods, which are widely used in machine learning.

We propose to use the nuclear norm of the exponential feature maps of the points as a regularizer in inverse problems, including denoising and recovery from incomplete data. We use the "kernel trick" to avoid the explicit computation of the maps or the surface, thus keeping the computational complexity manageable. We rely on an iterative reweighted algorithm to recover the denoised points. The resulting algorithm has similarities to iterative non-local methods [9, 10] that are widely used in image processing and graph signal processing. Specifically, it alternates between the estimation of a graph Laplacian, which specifies the connectivity of the points, and the smoothing of points guided by the graph Laplacian. The proposed framework can be thought of as a more principled alternative to the above heuristic approaches. We demonstrate the utility of the algorithm in the challenging application involving the recovery of free breathing and ungated cardiac MRI data from highly undersampled measurements. The experiments show the great potential of this method in imaging applications, where patches/images can be assumed to lie on a manifold or smooth surface in high-dimensional space.

This work is built upon our prior work [8] and the recent work by Ongie et al., which considered polynomial kernels [11]. The direct extension of [8] to recover the patch/image surface is computationally prohibitive due to the curse of dimensionality. The main focus of this work is to generalize [11] to shift invariant kernels, which are more widely used in practice. Our approach uses an implicit representation of the surface using bandlimited or Gaussian functions, which is more useful in practical applications. In addition, our work shows the connections with graph Laplacian based methods that are widely used in graph signal processing as well as patch-based methods. The application to MRI has similarities to [4, 6], which estimate the kernel subspace structure from navigator data. In contrast, our focus is to develop a navigator-free reconstruction strategy. Besides, [6] relies on an explicit mapping between the data and higher dimensional maps, whereas our scheme only works in the original domain.

2. BANDLIMITED SURFACES & ANNIHILATION

We assume the point cloud to be supported on a surface in $[-1/2, 1/2]^n$, which is the zero-level set of a bandlimited potential function:

$$\{\mathbf{r} \in \mathbb{R}^n \mid \psi(\mathbf{r}) = 0\} \text{ where } \psi(\mathbf{r}) = \sum_{\mathbf{k} \in \Lambda} \mathbf{c}_{\mathbf{k}} e^{j 2\pi \mathbf{k}^T \mathbf{r}} \quad (1)$$

Here, $\mathbf{c}_{\mathbf{k}}$; $\mathbf{k} \in \Lambda$ is the smallest set of coefficients (minimal set) that satisfies the above relation. $\Lambda \subset \mathbb{Z}^n$ is a set of contiguous locations that indicates the support of the Fourier series coefficients of ψ . Note that the above level set representation is widely used in image processing applications and can represent a large class of open and closed shapes [8].

Consider an arbitrary point \mathbf{x} on the above surface (1). By definition (1), we have the annihilation relation $\psi(\mathbf{x}) = \sum_{\mathbf{k} \in \Lambda} \mathbf{c}_{\mathbf{k}} e^{j2\pi\mathbf{k}^T \mathbf{x}} = 0$. The annihilation relation can be re-expressed in terms of the non-linear feature map $\phi_{\Lambda}(\mathbf{x})$ as $\mathbf{c}^T \phi_{\Lambda}(\mathbf{x}) = 0$, where $\phi: \mathbb{R}^n \rightarrow \mathbb{C}^{|\Lambda|}$ is defined as:

$$\phi_{\Lambda}(\mathbf{x}) = \left[e^{j2\pi\mathbf{k}_1^T \mathbf{x}} \dots e^{j2\pi\mathbf{k}_{|\Lambda|}^T \mathbf{x}} \right]^T \quad (2)$$

We now consider a collection of $N > |\Lambda|$ points on the surface, stacked into a matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]$. With the same argument as above, we have:

$$\mathbf{c}^T \underbrace{[\phi_{\Lambda}(\mathbf{x}_1), \dots, \phi_{\Lambda}(\mathbf{x}_N)]}_{\Phi_{\Lambda}(\mathbf{X})} = 0, \quad (3)$$

where $\Phi_{\Lambda}(\mathbf{X}) \in \mathbb{C}^{|\Lambda| \times N}$. The relation (3) implies that there is one vector in the null space of the feature matrix $\Phi_{\Lambda}(\mathbf{X})$. Since \mathbf{c} is the unique minimal filter, there is no other null space vector that satisfies the above relation. This implies that the rank of the feature matrix is $|\Lambda| - 1$. In practice, we do not know the exact Fourier support of ψ ; we can overestimate the support to $\Gamma \supseteq \Lambda$. If we considered a mapping $\phi_{\Gamma}: \mathbb{R}^n \rightarrow \mathbb{C}^{|\Gamma|}$, any function ψ' that is bandlimited to Γ is of the form $\psi'(\mathbf{r}) = \psi(\mathbf{r})\eta(\mathbf{r})$, where $\eta(\mathbf{r})$ is any function. Then, for any \mathbf{r} such that $\psi(\mathbf{r}) = 0$, we also have $\psi'(\mathbf{r}) = 0$. One can find $\Gamma: \Lambda$ linearly independent η functions, where $\Gamma: \Lambda$ denotes the set of translations of Λ in Γ [8]. We have:

$$\text{rank}(\Phi_{\Gamma}(\mathbf{X})) \leq |\Gamma| - |\Gamma: \Lambda| \quad (4)$$

provided $\Lambda \subseteq \Gamma$ and the points $\{\mathbf{x}_j\}$ are on the surface (1).

2.1. Surface recovery from noisy point cloud

The least square estimation of the coefficients from the data points can be posed as a minimization of the criterion $\mathcal{E}(\mathbf{c}) = \sum_{i=1}^N \|\psi(\mathbf{x}_i)\|^2 = \mathbf{c}^T \mathbf{Q}_{\Gamma} \mathbf{c}$, where

$\mathbf{Q}_{\Gamma} = \sum_{i=1}^N \phi_{\Gamma}(\mathbf{x}_i) \phi_{\Gamma}(\mathbf{x}_i)^T$. The estimation of the coefficients can be posed as the eigenvalue problem:

$$\mathbf{c}^* = \underset{\mathbf{c}}{\text{argmin}} \mathbf{c}^T \mathbf{Q}_{\Gamma} \mathbf{c} \text{ such that } \|\mathbf{c}\|^2 = 1, \quad (5)$$

whose solution is the minimum eigenvector of the matrix \mathbf{Q}_{Γ} .

Note that bandlimited ψ functions are often oscillatory, resulting in spurious zero locations. One may instead use weighted basis functions $\phi'(\mathbf{x}) = \mathbf{D} \phi(\mathbf{x})$ to obtain smoother curves. We choose the weight matrix as a diagonal matrix with diagonal entries $e^{-\pi^2 \sigma^2 \|\mathbf{k}\|^2}$.

2.2. Relation to shift invariant non-linear kernels

The above explicit approach is feasible when the dimension of the space n is small. However, the dimension of the feature space grows as a power of n , making this approach impractical in applications involving clouds of images or patches. Hence, we rely on the right null space relations. Since the rank of the feature matrix Φ_Γ is $r = |\Gamma| - |\Gamma : \Lambda|$, we can find $N - r$ vectors \mathbf{v}_i such that $\Phi_\Gamma(\mathbf{X}) \mathbf{v}_i = \mathbf{0}$, or equivalently,

$$\underbrace{\Phi_\Gamma(\mathbf{X})^H \Phi_\Gamma(\mathbf{X})}_{\mathbf{K}^\Gamma} \mathbf{v}_i = \mathbf{0}, \quad (6)$$

The entries of the $N \times N$ Gram matrix \mathbf{K}^Γ are:

$$\mathbf{K}_{i,j}^\Gamma = \phi_\Gamma(\mathbf{x}_i)^H \phi_\Gamma(\mathbf{x}_j) = \sum_{\mathbf{k} \in \Gamma} \underbrace{e^{j 2\pi \mathbf{k}^T (\mathbf{x}_j - \mathbf{x}_i)}}_{\kappa_\Gamma(\mathbf{x}_j - \mathbf{x}_i)}, \quad (7)$$

where $\kappa_\Gamma(\mathbf{r})$ is shift invariant. When Γ is a centered cube in \mathbf{R}^n , we have $\kappa_\Gamma(\mathbf{r}) = D_\Gamma(\mathbf{r})$, where D_Γ is the Dirichlet kernel whose shape is controlled by the support set Γ . The above arguments show that:

$$\text{rank}(\mathbf{K}_\Gamma) \leq |\Gamma| - |\Gamma : \Lambda| \quad (8)$$

provided the points $\mathbf{x}_i, i = 1, \dots, N$ lie on the surface (1).

If we choose the weighted maps $\phi' = \mathbf{D} \phi$, the kernel function approaches a Gaussian function, periodized to $[-1/2, 1/2]^n$ as $\Lambda \rightarrow \mathbb{Z}^n$. As $|\Gamma| \rightarrow \infty$, the matrix \mathbf{K}_Γ is theoretically full rank. However, we observe that the Fourier series coefficients of a Gaussian function can be safely approximated to be zero outside $|\mathbf{k}| < 3/\pi\sigma$, which translates to $|\Lambda| \approx \left(\frac{6}{\pi\sigma}\right)^n$; i.e., the rank will be small for high values of σ . We choose Gaussian kernels since they are more isotropic and less oscillatory than the Dirichlet kernel.

2.3. Recovery of noisy point clouds in high dimensions

We rely on the low rank structure of the kernel matrix \mathbf{K} to recover the noisy points. Specifically, with the addition of noise, the points deviate from the zero-level set of ψ . A high bandwidth potential function is needed to represent the noisy surface, which translates to a high rank matrix \mathbf{K} . Hence, we propose to use the nuclear norm of the feature vectors as a regularizer in the recovery of the cloud from noisy data:

$$\min_{\mathbf{X}} \|\mathcal{A}(\mathbf{X}) - \mathbf{b}\|^2 + \lambda \|\Phi(\mathbf{X})\|_* \quad (9)$$

For the case of recovery of points from noisy samples, $\mathcal{A}(\mathbf{X}) = \mathbf{X}$, while in the more general case \mathcal{A} could be an under-sampling operator. The above formulation is illustrated in Fig. 2, where we demonstrate the denoising of a shape from its samples; see caption for details. While we illustrate this approach in 2D, it is general enough to be applied in any dimension.

However, the direct evaluation of the maps is computationally prohibitive in higher dimensions. We hence now propose an efficient algorithm to denoise the points exploiting the low rank structure of the maps.

2.4. Iterative reweighted least squares (IRLS) algorithm

We use the IRLS algorithm to solve optimization problem (9). By the definition of the nuclear norm:

$$\|\Phi(\mathbf{X})\|_* = \text{trace}[(\Phi(\mathbf{X})^T \Phi(\mathbf{X}))^{\frac{1}{2}}] \approx \text{trace}[(\mathbf{X})\mathbf{Q}]$$

where $\mathbf{Q} = [\mathcal{K}(\mathbf{X}) + \gamma\mathbf{I}]^{-\frac{1}{2}}$. We use this property to solve (9) using an alternating strategy:

$$\mathbf{X}^{(n)} = \underset{\mathbf{X}}{\text{argmin}} \|\mathcal{A}(\mathbf{X}) - \mathbf{b}\|^2 + \lambda \text{trace}[\mathcal{K}(\mathbf{X})\mathbf{Q}^{(n-1)}] \quad (10)$$

where

$$\mathbf{Q}^{(n)} = [\mathcal{K}(\mathbf{X}^{(n)}) + \gamma^{(n)}\mathbf{I}]^{-\frac{1}{2}} \quad (11)$$

Note that the solution for (10) involves a system of non-linear equations. We use gradient linearization to simplify our computations. Since $\mathcal{K}(\mathbf{X})$ is a Gaussian kernel matrix, linearizing the gradient of the objective function of (10) at each iterate results in the following equivalent optimization problem for (10):

$$\mathbf{X}^{(n)} = \underset{\mathbf{X}}{\text{argmin}} \|\mathcal{A}(\mathbf{X}) - \mathbf{b}\|^2 + \lambda \text{trace}(\mathbf{X}^T \mathbf{L}^{(n-1)} \mathbf{X}), \quad (12)$$

where $\mathbf{L}^{(n)} = \mathbf{D}^{(n)} - \mathbf{W}^{(n)}$. Here, the entries of $\mathbf{W}^{(n)}$ are $\mathbf{W}_{ij}^{(n)} = \frac{f'(\mathbf{x}_i^{(n)} - \mathbf{x}_j^{(n)})}{\|\mathbf{x}_i^{(n)} - \mathbf{x}_j^{(n)}\|} \mathbf{Q}_{ij}^{(n)}$ and

$\mathbf{D}_{ii}^{(n)} = \sum_j \mathbf{W}_{ij}^{(n)}$. The function f determining the weight matrix depends on the Radial Basis Function kernel as: $[\mathcal{K}(\mathbf{X})]_{ij} = f(\|\mathbf{x}_i - \mathbf{x}_j\|)$.

We observe the equivalence of the above optimization strategy with widely used non-local means and graph optimization schemes. These schemes estimate a Laplacian matrix \mathbf{L} or equivalently a weight matrix \mathbf{W} , followed by the minimization of the cost function (12). These approaches can thus be seen as fitting a smooth bandlimited surface to the point cloud of patches or signals that are assumed to be on the graph.

For the denoising problem (where $\mathcal{A}(\mathbf{X}) = \mathbf{X}$), and for some operators \mathcal{A} it is convenient to solve problem (12) analytically. In other cases, a conjugate-gradient algorithm can be used to solve it. Note that in the proposed IRLS iterations, it is sufficient to only compute the kernel matrix $\mathcal{K}(\mathbf{X})$; the matrix of feature vectors $\Phi(\mathbf{X})$ is never required to be computed explicitly. This is important since the feature vectors in many cases may be large or even infinite dimensional.

2.5. Application to cardiac MRI

We apply the proposed framework to the recovery of free breathing and ungated MRI data from highly undersampled measurements. Since MRI is a slow imaging modality, the standard practice in functional cardiac MRI is to integrate the data from multiple heart beats in a breath-held acquisition. However, this approach is often challenging in pediatric and obese subjects who cannot hold their breath. The image frames can be safely assumed to be non-linear functions of two parameters: cardiac and respiratory phase, and hence can be modeled as points on a smooth manifold in high dimensional space. In our previous work, we have acquired the data with navigators, from which the graph Laplacian is estimated. The main challenge with this strategy is the need for customized sequences and the additional 40% overhead in acquiring the navigators. In this paper we enable a navigator-free acquisition scheme using the proposed approach.

3. RESULTS

We use the algorithm (9) to directly recover the dynamic MRI dataset \mathbf{X} from golden angle acquired dataset with 10 radial spokes per frame (50 fold acceleration) and a 30 channel cardiac and spine array. The data was acquired in 40 seconds, which corresponds to 1000 frames and a temporal resolution of ≈ 40 ms. A subset of 15 coils was chosen for the reconstruction. The proposed technique was used to iteratively estimate the \mathbf{L} matrix from the 31×31 centre under-sampled k-space data. This \mathbf{L} matrix was used to reconstruct the full data in a single iteration. The in-vivo cardiac free-breathing images recovered from their ≈ 50 fold undersampled measurements using the proposed scheme are shown in Fig. 3. Two reconstructed frames are shown from the dataset along with the temporal profile. It is observed that the proposed scheme can reconstruct images of good quality while preserving the temporal dynamics.

4. CONCLUSION

We introduce a novel framework for the recovery of images/patches, which lie on a smooth low-dimensional manifold/surface in high dimensional space. We model the surface as the level set of a bandlimited function. We show that the non-linear feature maps of the points satisfy annihilation relations. Since these relations imply that the maps lie in a finite dimensional subspace, we use the nuclear norm of the maps as a regularizer to recover the points/images from highly undersampled measurements. The application of this framework to navigator-free free breathing and ungated cardiac MRI provides promising results.

Acknowledgments

This work is supported by NIH 1R01EB019961-01A1 and onrn000141310202.

5. REFERENCES

- [1]. Danielyan A, Katkovnik V, and Egiazarian K, "Bm3d frames and variational image deblurring," IEEE Tran. Image Processing, vol. 21, no. 4, pp. 1715–1728, 4 2012.

- [2]. Mohsin Yasir Q., Lingala Sajan Goud, DiBella Edward, and Jacob Mathews, "Accelerated dynamic mri using patch regularization for implicit motion compensation," *Magnetic Resonance in Medicine*, vol. 77, no. 3, pp. 1238–1248, 2017. [PubMed: 27091812]
- [3]. Mohsin Yasir Q, Ongie Gregory, and Jacob Mathews, "Iterative shrinkage algorithm for patch-smoothness regularized medical image recovery," *IEEE transactions on medical imaging*, vol. 34, no. 12, pp. 2417–2428, 2015. [PubMed: 25647445]
- [4]. Poddar S and Jacob M, "Dynamic mri using smoothness regularization on manifolds (storm)," *IEEE Tran. Medical Imaging*, vol. 35, no. 4, pp. 1106–1115, 4 2016.
- [5]. Poddar Sunrita, Lingala Sajan Goud, and Jacob Mathews, "Joint recovery of under sampled signals on a manifold: Application to free breathing cardiac mri," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on IEEE*, 2014, pp. 6904–6908.
- [6]. Nakarmi U, Wang Y, Lyu J, Liang D, and Ying L, "A kernel-based low-rank (klr) model for low-dimensional manifold recovery in highly accelerated dynamic mri," *IEEE Transactions on Medical Imaging*, vol. PP, no. 99, pp. 1–1, 2017.
- [7]. Bhatia Kanwal K., Caballero Jose, Price Anthony N., Sun Ying, Hajnal Jo V., and Rueckert Daniel, *Fast Reconstruction of Accelerated Dynamic MRI Using Manifold Kernel Regression*, pp. 510–518, Springer International Publishing, Cham, 2015.
- [8]. Ongie Greg and Jacob Mathews, "Off-the-grid recovery of piecewise constant images from few fourier samples," *SIAM Journal on Imaging Sciences*, vol. 9, no. 3, pp. 1004–1041, 2016. [PubMed: 29973971]
- [9]. Yang Z and Jacob M, "Nonlocal regularization of inverse problems: A unified variational framework," *IEEE Transactions on Image Processing*, vol. 22, no. 8, pp. 3192–3203, 8 2013. [PubMed: 23014745]
- [10]. Kalofolias Vassilis, "How to learn a graph from smooth signals," in *Proc. Int. Conf. Artificial Intelligence and Statistics*, 2016, vol. 51, pp. 920–929.
- [11]. Ongie Greg, Willett Rebecca, Nowak Robert D, and Balzano Laura, "Algebraic variety models for high-rank matrix completion," *arXiv preprint arXiv:1703.09631*, 2017.

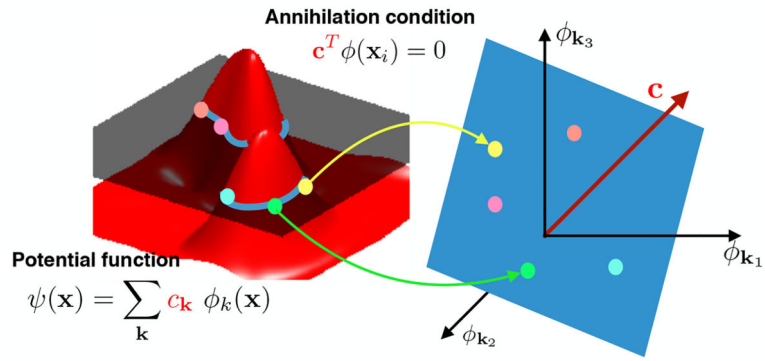




Fig. 1: Illustration of the annihilation relations in 2-D. We assume that the curves/surface is the zero-level set of the potential function $\psi(\mathbf{x}) = 0 = \mathbf{c}^T \phi(\mathbf{x}_i)$, which can be seen as an annihilation relation in the non-linear feature space $\phi(\mathbf{x})$. Specifically, the maps of the points lie on a plane orthogonal to \mathbf{c} .



(a) Original & Bandlimited (b) Noisy & Denoised



(c) Singular value decay



(d) Potential function

Fig. 2:

Recovery of shape from noisy samples of the Tigerhawk logo. (a) shows the original points, and their projection to a bandlimited Gaussian curve with $\sigma = 0.15$. Note from (c) that the singular values of the bandlimited points decay rapidly, while the ones corresponding to the noisy points decay slower. The denoising algorithm (9) exploits the lower nuclear norm of $\Phi(\mathbf{X})$ to denoise the data. The points after random noise addition and denoising using our algorithm is displayed in (b). Since $\Phi(\mathbf{X})$ has several null space vectors, we display the sum of squares function in (d).

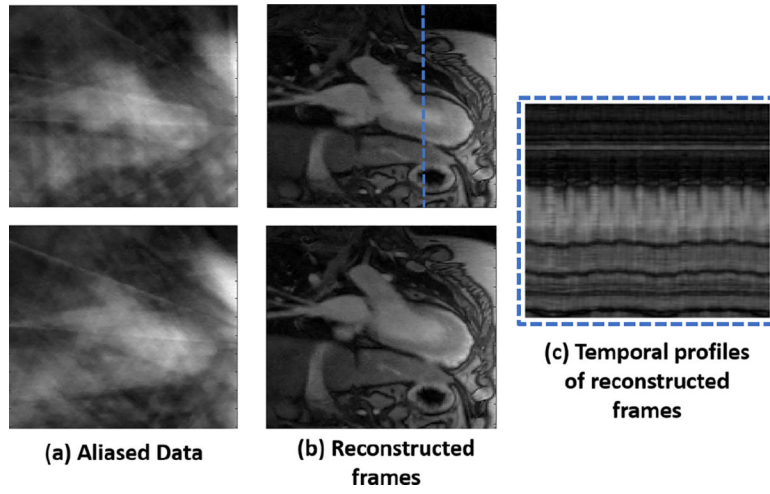


Fig. 3: Reconstruction of free breathing and ungated cardiac MRI data from ≈ 50 fold undersampled measurements. Recovery of the data using standard CS or low-rank methods is challenging due to the extensive motion; see [4] for comparisons with state of the art methods. (a) denotes two frames, recovered by gridding. (b) Corresponding frames reconstructed using the proposed scheme. (c) Temporal profile from reconstructed frames at the position marked by dotted blue line. The image quality of the proposed scheme is comparable to breath-held methods.