Sparse Signal Recovery via Correlated Degradation Model

Nasser Eslahi[⋆], Venkatessh Ramakrishnan[†], Kaj Wiik[‡] and Alessandro Foi[⋆]
[⋆]Department of Signal Processing, Tampere University of Technology, Finland
[†]Departmento de Astronomia, Universidad de Concepción, Chile
[‡]Tuorla Observatory, University of Turku, Finland

Sparse signal recovery aims to recover an unknown signal $\mathbf{x}\in\mathbb{R}^n$ from few non-adaptive, possibly noisy, linear measurements $\mathbf{y}\in\mathbb{R}^m$ using a nonlinear sparsity-promoting algorithm, under the assumption that $\mathbf{x}$ is sparse or compressible with respect to a known basis or frame [1]. Specifically, $\mathbf{y}=\mathbf{A}\mathbf{x}+\mathbf{e}$, where $\mathbf{A}\in\mathbb{R}^{m\times n}$ is the measurement matrix, $\mathbf{e}\in\mathbb{R}^m$ is the measurement error, and $m\ll n$. Many of the sparse recovery approaches proposed during the past decade can be regarded as iterative estimation of a signal from a degraded observation [2]–[6]. A common feature of these approaches is that the degradation is modeled as independent and identically distributed (i.i.d.) additive noise that has to be alleviated by a denoising filter at each iteration. For instance, Approximate Message Passing (AMP) [2] enforces an i.i.d. additive Gaussian noise denoising problem at each iteration through the Onsager correction term [2]–[4]. *Plug&Play*-Prior (P^3) frameworks [5], [6] use denoising algorithms as regularizers (priors) for model-based inversion via the alternating direction method of multipliers (ADMM) [7]. The idea of ADMM is to convert an unconstrained optimization problem $\hat{\mathbf{x}}=\operatorname{argmin}_{\mathbf{x}} f(\mathbf{x})+\lambda g(\mathbf{x})$ into its equivalent constrained form which is then decoupled into two separate optimizations:

$$(\mathbf{x}^{k+1},\mathbf{v}^{k+1})=\operatorname*{argmin}_{\mathbf{x},\mathbf{v}\in\mathbb{R}^n}\left\{f(\mathbf{x})+\lambda g(\mathbf{v})+\frac{\mu}{2}\|\mathbf{x}-\mathbf{v}-\mathbf{b}\|_2^2\right\}, \quad (1)$$
$$\mathbf{b}^{k+1}=\mathbf{b}^{k}-(\mathbf{x}^{k+1}-\mathbf{v}^{k+1}),$$

where $\mathbf{v}$ is the auxiliary variable, $\mathbf{b}$ is the scaled Lagrange multiplier, and μ is a fixed parameter for improving the numerical stability of the algorithm. Note that the $\mathbf{x}$-subproblem and $\mathbf{v}$-subproblem of (1) are the Moreau proximity operators of f and g computed at $\mathbf{v}^k+\mathbf{b}^k$ and $\mathbf{x}^{k+1}-\mathbf{b}^k$, respectively. P^3 suggests that the $\mathbf{x}$-subproblem can be interpreted as an inversion step since it only depends on the choice of forward model $f(\mathbf{x})$, while the $\mathbf{v}$-subproblem can be regarded as a denoising step since it only depends on the choice of prior $g(\mathbf{v})$. Thus, one can reformulate the $\mathbf{v}$-subproblem of (1) as $\mathbf{v}^{k+1}=\mathcal{D}(\tilde{\mathbf{v}}^k)\stackrel{\text{def}}{=}\operatorname{argmin}_{\mathbf{v}} g(\mathbf{v})+\frac{1}{2\sigma^2}\|\mathbf{v}-\tilde{\mathbf{v}}^k\|_2^2$, where $\tilde{\mathbf{v}}^k=\mathbf{x}^{k+1}-\mathbf{b}^k$ is treated as some noisy (degraded) realization of an unknown signal $\mathbf{v}$, and $\mathcal{D}:\mathbb{R}^n\to\mathbb{R}^n$ is a generic denoiser. Various denoisers can be applied as priors in P^3, yet, to the best of our knowledge, all works have so far adopted filters modeling the degradation as i.i.d. noise.

In this work, we focus on the modeling of such degradations as spatially (or spatiotemporally) correlated noise. Our motivation is that the common assumption of i.i.d. Gaussian noise is valid only under special conditions that are hardly met in practice, e.g., like the measurement matrix $\mathbf{A}$ being itself random i.i.d. Gaussian (as required by AMP). In fact, in many practical sparse signal recovery applications, there can be a significant spatial (and temporal) correlation in the noise $\mathbf{b}^k$ at each iteration (see Figures 1 and 2). In contrast to i.i.d. white noise, correlated noise can lead to disproportions in the magnitude of errors across the data spectrum, to an extent that i.i.d. denoisers may not effectively discern between the true signal and noise in regularization via shrinkage. Hence, ignoring such correlation in the denoising step can lead to ineffective filtering and also distortion to the underlying signal.

Our contribution is to model the noise $\mathbf{b}^k$ at each iteration as spatially (or spatiotemporally) correlated noise, and adopt such model within the denoising. Specifically, we consider the noise power spectral densities (PSD) with respect to the sparsifying transforms used internally by the filters: these PSDs modulate the shrinkage thresholds, i.e. allow to compare the magnitude of each transform coefficient against that of the corrupting noise. While more sophisticated than an i.i.d. model, our correlated noise modeling remains in general only an approximation of the true degradation statistics at each iteration, which can be intricate and are result of multiple contributors: the structure of $\mathbf{A}$, the statistics of $\mathbf{e}$, as well as their interaction with the structure of $\mathbf{x}$ and the effect of $\mathcal{D}$ during the previous iterations. Therefore, we pragmatically resort to fitting a PSD description to the residuals $\mathbf{b}^k$ via the median absolute deviation (MAD) of the transform spectra of all patches. This is performed at every iteration and constitutes an adaptive estimate of the noise correlation which can be utilized by the filter $\mathcal{D}$ in the denoising step.

Experiments were performed on three different applications of sparse recovery: volumetric magnetic resonance imaging (MRI), multi-epoch radio-interferometry (RI), and video recovery from temporally compressed measurements. The $64\times64\times64$-voxel volumetric MRI reconstruction of the BrainWeb phantom [8] with nonzero phase, subsampling ratio of 30% with spherical sampling trajectory and initial noise $\sigma_{\mathbf{e}}=5\%$, is shown in Figure 3. We use the BM4D filter [9] as denoiser for volumetric data, under the modeling assumption of either i.i.d. (as in [9]) or spatially correlated noise (this work). BM4D operates by grouping mutually similar 3-D cubes of voxels into 4-D arrays; the 4-D PSD of each group is obtained by replicating the input 3-D PSD of the cubes (see Fig. 1) under the simplifying assumption that grouped cubes do not overlap. Multi-epoch RI images and spatiotemporally encoded video are recovered via P^3 using the RF3D denoiser [10] for regularization (Figures 4 and 5). RF3D operates on spatiotemporal volumes formed by concatenating blocks along motion trajectories. The PSD of each spatiotemporal volume is computed internally by the filter, based on the volume trajectory, from a pair of 2D PSDs, one of the temporally uncorrelated noise and another of the fixed-pattern noise component.

The results demonstrate that modeling the degradation as correlated noise yields a better detection and sharper recovery of fine structures, attesting its practical advantage over the common i.i.d. modeling in iterative sparse signal recovery. The PSDs shown in Figures 1 and 2 give evidence of order-of-magnitude differences across the spectrum, with noticeable anisotropy. While here we report experiments with noisy measurements, we note that similar forms of correlation take place also in the sparse recovery from noise-free (i.e. $\mathbf{e}=0$) measurements.

The research leading to these results has received funding from the European Union's H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement no. 642685 MacSeNet, and by the Academy of Finland (project no. 252547).

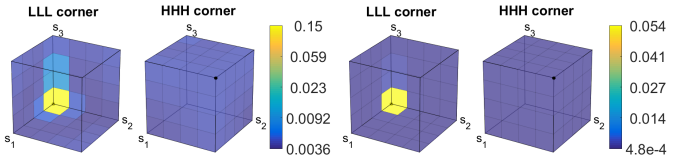

Figure 1. Spatial correlation in the volumetric residual of MRI reconstruction (see Fig. 3) at iteration 100 (left) and 2000 (right) of 2000, visualized as noise power spectrum (PSD) with respect to the 4×4×4 3D DCT. Only the lowest-frequency and highest-frequency faces of each 3D PSD cube are shown.

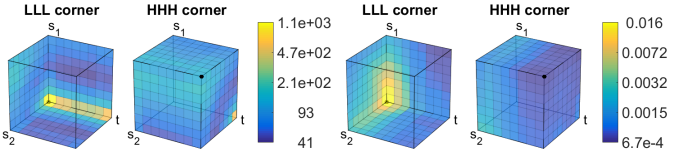

Figure 2. Spatiotemporal correlation in the residual in multi-epoch RI imaging (left, iter. 30/70, see Fig. 4) and video recovery (right, iter. 6/18, see Fig. 5), visualized as PSD with respect to the 8×8×8 3D DCT. Only the lowest-frequency and highest-frequency faces of each 3D PSD cube are shown.

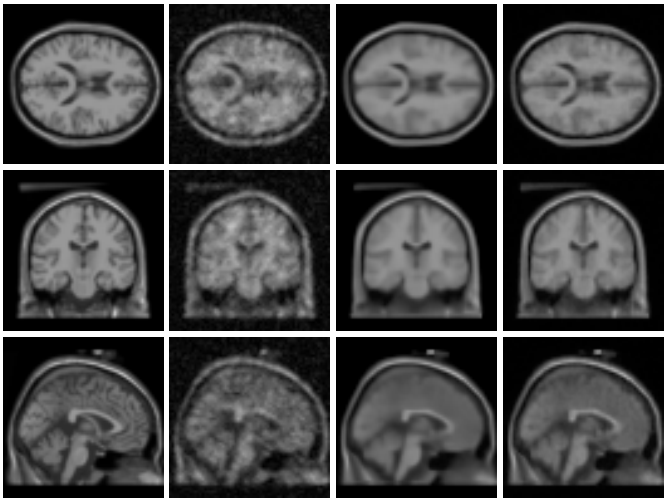

Figure 3. 1st (resp. 2nd and 3rd) row, from left to right, 2D transverse (resp. coronal and sagittal) cross-sections of: noise-free BrainWeb phantom; initial back projections (PSNR 23.85 dB); and final estimates of the magnitude via i.i.d. (27.86 dB) and correlated degradation modeling (28.93 dB), respectively. PSNR is for magnitude only, as the most informative component.

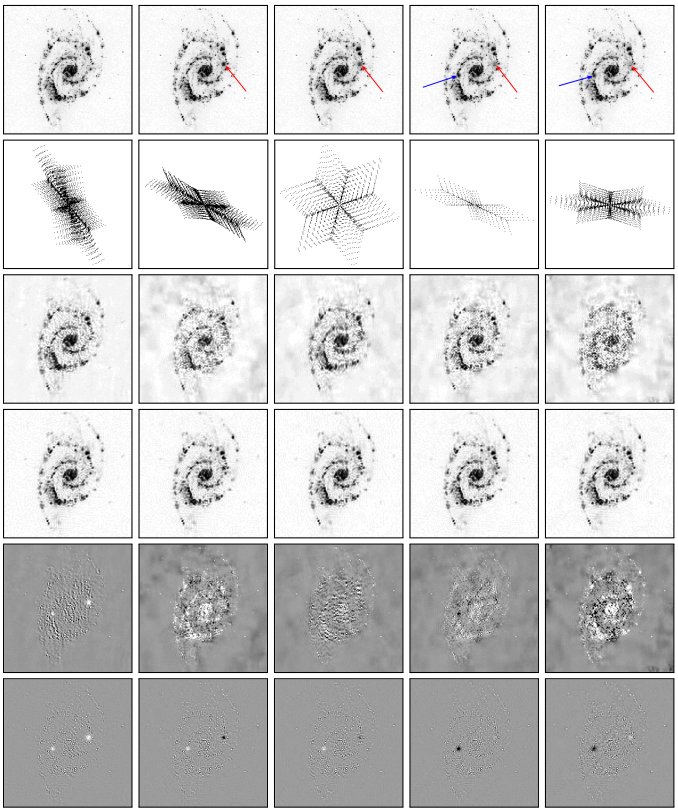

Figure 4. 1st row: ground-truth galaxy images with two smoothly varying transient sources at different epochs – shown by red and blue arrows. 2nd row: visibilities coverage in the Fourier plane. 3rd and 4th rows: the recovered RI images via i.i.d. and correlated degradation model, respectively. 5th and 6th rows: subtraction of the recovered images of third and fourth rows from median of the ground truth images. Epochs #1,#3,#5,#7,#9 are shown from the total of 10 epochs in the dataset. The transient source indicated by the red arrow appears, fades and eventually disappears during epochs #2–#9, whereas the other one appears and fades during epochs #6–#10.

Figure 5. From left to right column: original frames (# 10 & 39) of the *Traffic* 256×256×48 video sequence; corresponding spatiotemporally encoded frames (each being the sum of 4 consecutive masked frames from 4 pseudorandom binary masks); the recovered frames using a spatiotemporal recovery approach via i.i.d. and correlated degradation modeling, respectively.

References

[1] E. J. Candès, J. K. Romberg, and T. Tao, "Stale signal recovery from incomplete and inaccurate measurements," *Comm. Pure and Applied Math.*, vol. 59, no. 8, pp. 1207-1223, 2006.

[2] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," *Proc. Nat. Acad. Sci.*, vol. 106, no. 45, pp. 18914-18919, 2009.

[3] C. A. Metzler, A. Maleki, and R. G. Baraniuk, "From denoising to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 62, no. 9, pp. 5117-5144, 2016.

[4] J. Tan, Y. Ma, and D. Baron, "Compressive imaging via approximate message passing with image denoising," *IEEE Trans. Signal Process.*, vol. 63, no. 8, pp. 2085-2092, 2015.

[5] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Global Conf. Signal and Inf. Process. (GlobSIP)*, 2013, pp. 945-948.

[6] S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-play ADMM for image restoration: fixed point convergence and applications," *IEEE Trans. Comput. Imag.*, to be published, doi: 10.1109/TCI.2016.2629286.

[7] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1-122, 2011.

[8] R. Vincent, "Brainweb: Simulated brain database," 2006. Available: http://mouldy.bic.mni.mcgill.ca/brainweb/

[9] M. Maggioni, V. Katkovnik, K. Egiazarian, and A. Foi, "Nonlocal transform-domain filter for volumetric data denoising and reconstruction," *IEEE Trans. Imag. Process.*, vol. 22, no. 1, pp. 119-133, 2013.

[10] M. Maggioni, E. Sanchez-Monge, and A. Foi, "Joint removal of random and fixed-pattern noise through spatiotemporal video filtering," *IEEE Trans. Imag. Process.*, vol. 23, no. 10, pp. 4282-4296, 2014.