

GROUP-SPARSE MATRIX RECOVERY

Xiangrong Zeng and Mário A. T. Figueiredo

Instituto de Telecomunicações, Instituto Superior Técnico, Universidade de Lisboa, Portugal

ABSTRACT

We apply the OSCAR (octagonal selection and clustering algorithms for regression) in recovering group-sparse matrices (two-dimensional—2D—arrays) from compressive measurements. We propose a 2D version of OSCAR (2OSCAR) consisting of the ℓ_1 norm and the pair-wise ℓ_∞ norm, which is convex but non-differentiable. We show that the proximity operator of 2OSCAR can be computed based on that of OSCAR. The 2OSCAR problem can thus be efficiently solved by state-of-the-art proximal splitting algorithms. Experiments on group-sparse 2D array recovery show that 2OSCAR regularization solved by the SpaRSA algorithm is the fastest choice, while the PADMM algorithm (with debiasing) yields the most accurate results.

Index Terms— group sparsity, matrix recovery, proximal splitting algorithms, proximity operator, signal recovery.

1. INTRODUCTION

The problem studied in this paper is the classical one of recovering \mathbf{X} from

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{W}, \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a known sensing matrix, $\mathbf{X} \in \mathbb{R}^{n \times d}$ the original unknown matrix/2D-array, $\mathbf{Y} \in \mathbb{R}^{m \times d}$ is the observed data, and $\mathbf{W} \in \mathbb{R}^{m \times d}$ denotes additive noise. In many cases of interest, we have $m < n$, making (1) an ill-posed problem, which can only be addressed by using some form of regularization that injects prior knowledge about the unknown \mathbf{X} . Classical regularization formulations seek solutions of problems of the form

$$\min_{\mathbf{X}} F(\mathbf{X}) + \Phi(\mathbf{X}), \quad (2)$$

or one of the equivalent (under mild conditions) forms

$$\min_{\mathbf{X}} \Phi(\mathbf{X}) \text{ s.t. } F(\mathbf{X}) \leq \varepsilon, \text{ or } \min_{\mathbf{X}} F(\mathbf{X}) \text{ s.t. } \Phi(\mathbf{X}) \leq \epsilon, \quad (3)$$

where $F(\mathbf{X})$ is the data-fidelity term and $\Phi(\mathbf{X})$ is the regularizer, the purpose of which is to enforce certain properties

on \mathbf{X} , such as sparsity or group sparsity, and ε and ϵ are non-negative parameters.

Problem (1) is more challenging than the widely studied linear inverse problem of recovering a vector $\mathbf{x} \in \mathbb{R}^n$ from

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w}, \quad (4)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$ is also a known sensing matrix, $\mathbf{y} \in \mathbb{R}^m$ is the observed vector, and $\mathbf{w} \in \mathbb{R}^m$ is additive noise. Comparing with (4), the matrix \mathbf{X} of interest in (1) is always assumed to be, not only sparse, but also to have a particular sparse structure. For instance, in the *multiple measurement vector* model [1], [2], [3], \mathbf{X} is an unknown source matrix that should be row sparse; in group LASSO [4], [5], [6], \mathbf{X} is a coefficient matrix that is also enforced to be row sparse; in *multi-task learning* [7], [8], \mathbf{X} is a task parameter matrix, which is usually assumed to be row or/and column sparse. In this paper, we pursue more general sparsity patterns for \mathbf{X} , that is, the arrangement of each group of nonzeros in \mathbf{X} is not limited to rows and/or columns, but may include row/columns segments, blocks, or other groups of connected non-zero elements. Before addressing the question of whether or not there are any available regularizers able to promote this kind of group sparsity, we first briefly review existing group-sparsity-inducing regularizers.

In recent years, much attention has been devoted not only to the sparsity of solutions, but also the structure of this sparsity [9]. In other words, not only the number of non-zeros in the solutions, but also how these non-zeros are located, are of interest. This research direction has lead to the concept of group/block sparsity [4], [10], or more general structured sparsity patterns [11], [12], [13]. A classical model for group sparsity is the *group LASSO* [4], which, making use of more information than the original LASSO [14] (namely, the structure of the groups) is able to simultaneously encourage sparsity and group sparsity. In addition, the *sparse group LASSO* approach, proposed in [15] uses a regularizer that consists of an ℓ_1 term plus the group LASSO regularizer, thus unlike group LASSO, it not only selects groups, but also individual variables within each group.

It has also been observed that in some real-world problems, it makes sense to encourage the solution, not only to be sparse, but also to have several components sharing similar values. To formalize this goal, several generic models have been proposed, such as the *elastic net* [16], the *fused LASSO*

Work partially supported by Fundação para a Ciência e Tecnologia, grants PEst-OE/EEI/LA0008/2013 and PTDC/EEI-PRO/1470/2012. Xiangrong Zeng is partially supported by grant SFRH/BD/75991/2011.




Fig. 1. Illustration of LASSO, elastic net, fused LASSO and OSCAR

[17], and the *octagonal shrinkage and clustering algorithm for regression* (OSCAR) [18].

The level curves of several of the regularizers mentioned in the previous paragraph (for the 2D case) are shown in Fig. 1. The figure illustrates why these models promote variable grouping (unlike LASSO). Firstly, the regularizer of the elastic net [16] consists of an ℓ_1 term and an ℓ_2 term, thus simultaneously promoting sparsity and group-sparsity, in which the former comes from the sparsity-inducing corners (see Fig. 1) while the latter from its strictly convex edges, which creates a grouping effect similar to a quadratic/ridge regularizer. Secondly, the regularizer of the fused LASSO is composed of a ℓ_1 term and a total variation term, which encourages successive variables (in a certain order) to be similar, making it able to promote both sparsity and smoothness. Thirdly, the OSCAR regularizer (proposed by Bondell and Reich [18]) includes an ℓ_1 term and a pair-wise ℓ_∞ term, where the latter promotes equality (in absolute value) between each pair of variables.

There are some recent variants of the above group-sparsity regularizers, such as the *weighted fused LASSO*, presented in [19]. The *pair-wise fused LASSO* [19], which uses the pair-wise term of OSCAR, extends the fused LASSO to cases where the variables have no natural ordering. A novel *graph-guided fused LASSO* was proposed in [20], where the grouping structure is modeled by a graph. A Bayesian version of the elastic net was developed in [21]. Finally, an adaptive grouping pursuit method was proposed in [22], but the underlying regularizer is neither sparsity-promoting nor convex.

The fused LASSO, elastic net, and OSCAR regularizers all have the ability to promote sparsity and variable grouping. However, as pointed out in [23], OSCAR outperforms the other two models in terms of grouping. Moreover, the fused LASSO is not suitable for group according to magnitude, and the grouping ability of the convex edges of the elastic net is inferior to that of OSCAR. Thus, this paper will focus on the OSCAR regularizer to solve the problems of group-sparse matrix recovery.

In this paper, we propose a two-dimensional (matrix) version of OSCAR (2OSCAR) for group-sparse matrix recovery. Solving OSCAR regularization problems has been addressed in our previous work [24], in which, six state-of-the-art proximal splitting algorithms: FISTA [25], TwIST [26], SpaRSA [27], ADMM [28], SBM [29] and PADMM [30] are investigated. Naturally, we build the relationship between OSCAR and 2OSCAR, and then address 2OSCAR regularization problems as in [24].

Terminology and Notation

We denote vectors or general variables by lower case letters, and matrices by upper case ones. The ℓ_1 norm of a vector $\mathbf{x} \in \mathbb{R}^n$ is $\|\mathbf{x}\|_1 = \sum_{i=1}^n |x_i|$ where x_i represents the i -th component of \mathbf{x} , and that of a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ is $\|\mathbf{X}\|_1 = \sum_{i=1}^n \sum_{j=1}^d |X_{(i,j)}|$ where $X_{(i,j)}$ the entry of \mathbf{X} at the i -th row and the j -th column. Let $\|\mathbf{X}\|_F = (\sum_{i=1}^n \sum_{j=1}^d X_{(i,j)}^2)^{1/2}$ be the Frobenius norm of \mathbf{X} .

We now briefly review a fundamental object of convex analysis, the proximity operator, which will be used below. Let \mathcal{H} be a real Hilbert space with inner product $\langle \cdot, \cdot \rangle$ and norm $\|\cdot\|$. Let $f : \mathcal{H} \rightarrow [-\infty, +\infty]$ be a function and Γ be the class of all lower semi-continuous, convex, proper functions (not equal to $+\infty$ everywhere and never equal to $-\infty$). The proximity operator [31] of λf (where $f \in \Gamma$ and $\lambda \in \mathbb{R}_+$) is defined as

$$\text{prox}_{\lambda f}(\mathbf{v}) = \arg \min_{\mathbf{x} \in \mathcal{H}} \left(\lambda f(\mathbf{x}) + \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|^2 \right). \quad (5)$$

2. RECOVERING GROUP-SPARSE MATRICES

2.1. OSCAR and its 2D Version (2OSCAR)

The OSCAR criterion is given by [18]

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda_1 \|\mathbf{x}\|_1 + \lambda_2 \underbrace{\sum_{i < j} \max\{|x_i|, |x_j|\}}_{\Phi_{\text{OSCAR}}(\mathbf{x})}, \quad (6)$$

where the ℓ_2 term seeks data-fidelity, while the regularizer $\Phi_{\text{OSCAR}}(\mathbf{x})$ consists of an ℓ_1 term (promoting sparsity) and a pair-wise ℓ_∞ term ($((n(n-1))/2)$ pairs in total) encouraging equality (in magnitude) of each pair of elements $|x_i|$ and $|x_j|$. Thus, $\Phi_{\text{OSCAR}}(\mathbf{x})$ promotes both sparsity and grouping. Parameters λ_1 and λ_2 are nonnegative constants controlling the relative weights of the two terms. If $\lambda_2 = 0$, (6) becomes the LASSO, while if $\lambda_1 = 0$, Φ_{OSCAR} becomes a pair-wise ℓ_∞ regularizer. Note that, for any choice of $\lambda_1, \lambda_2 \in \mathbb{R}_+$, $\Phi_{\text{OSCAR}}(\mathbf{x})$ is convex and its ball is octagonal in the 2-D case. In the 2D case, the 8 vertices of this octagon can be divided into two categories: four sparsity-inducing vertices (located on the axes) and four equality-inducing vertices (see Fig. 1). Fig. 2 depicts

the a data-fidelity term and $\Phi_{\text{OSCAR}}(\mathbf{x})$, illustrating its possible effects.




Fig. 2. Illustration of the ℓ_2 term and $\Phi_{\text{OSCAR}}(\mathbf{x})$ in the $n = 2$ case. In this example, for the same least square solution $(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{y}$, the high correlation contour is more likely to hit the equality-inducing (grouping) vertex, whereas the low correlation contour “prefers” the sparsity-inducing vertex.

As discussed in Section 1, compared with group LASSO, OSCAR doesn’t require a pre-specification of group structure; compared with fused LASSO, it doesn’t depend on a certain order of the variables; compared with the elastic net, it has a stronger equality-inducing ability. All these features make OSCAR a convenient regularizer in many applications. A fundamental building block for using OSCAR is its proximity operator, $\text{prox}_{\Phi_{\text{OSCAR}}}$, which can be obtained exactly or approximately by the *grouping proximity operator* and the *approximate proximity operator* [24], respectively.

To address the inverse problem (1), we propose a matrix version of OSCAR, termed 2OSCAR, given by

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times d}} \frac{1}{2} \|\mathbf{Y} - \mathbf{A}\mathbf{X}\|_F^2 + \Phi_{2\text{OSCAR}}(\mathbf{X}) \quad (7)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{Y} \in \mathbb{R}^{m \times d}$ and

$$\Phi_{2\text{OSCAR}}(\mathbf{X}) = \Phi_{\text{OSCAR}}(\text{vec}(\mathbf{X})) \quad (8)$$

with vec denoting the vectorization function, which transforms a matrix into a vector by stacking the columns on top of each other.

Observing that, for any matrix \mathbf{Z} , we have $\|\mathbf{Z}\|_F^2 = \|\text{vec}(\mathbf{Z})\|_2^2$, we can write proximity operator of 2OSCAR as

$$\text{prox}_{\Phi_{2\text{OSCAR}}}(\mathbf{Z}) = \text{vec}^{-1}(\text{prox}_{\Phi_{\text{OSCAR}}}(\text{vec}(\mathbf{Z}))). \quad (9)$$

where $\text{prox}_{\Phi_{\text{OSCAR}}}$ can be obtained by the algorithm proposed in [24] and vec^{-1} is the inverse of the vectorization function, that is, it takes a nd vector and yields an $n \times d$ matrix.

2.2. Algorithms

In the 2OSCAR problem in (7), it is clear that the objective function is convex (since both terms are convex) and coercive (that is, it goes to ∞ as $\|\mathbf{X}\| \rightarrow \infty$), thus the set of minimizers is not empty.

To solve (7), we investigate six state-of-the-art proximal splitting algorithms: FISTA [25], TwIST [26], SpaRSA [27], ADMM [28], SBM [29] and PADMM [30]. Due to limitation of space, we next only detail SpaRSA, since it has been experimentally shown to be the fastest one. However, we will report below experimental results with the aforementioned six algorithms.

SpaRSA [27] is a fast proximal splitting that uses the step-length selection method of Barzilai and Borwein [32]. Its application to solve 2OSCAR problems leads to the following algorithm:

Algorithm SpaRSA for 2OSCAR

1. Set $k = 1$, $\eta > 1$, $\alpha_0 = \alpha_{\min} > 0$, $\alpha_{\max} > \alpha_{\min}$, and \mathbf{X}_0 .
2. $\mathbf{V}_0 = \mathbf{X}_0 - \mathbf{A}^T(\mathbf{A}\mathbf{X}_0 - \mathbf{Y})/\alpha_0$
3. $\mathbf{X}_1 = \text{Prox}_{\Phi_{2\text{OSCAR}}/\alpha_0}(\mathbf{V}_0)$
4. **repeat**
5. $\mathbf{S}_k = \mathbf{X}_k - \mathbf{X}_{k-1}$
6. $\mathbf{R}_k = \mathbf{A}^T \mathbf{A} \mathbf{S}_k$
7. $\hat{\alpha}_k = \frac{(\mathbf{S}_k)^T \mathbf{R}_k}{(\mathbf{S}_k)^T \mathbf{S}_k}$
8. $\alpha_k = \max\{\alpha_{\min}, \min\{\hat{\alpha}_k, \alpha_{\max}\}\}$
9. **repeat**
10. $\mathbf{V}_k = \mathbf{X}_k - \mathbf{A}^T(\mathbf{A}\mathbf{X}_k - \mathbf{Y})/\alpha_k$
11. $\mathbf{X}_{k+1} = \text{Prox}_{\Phi_{2\text{OSCAR}}/\alpha_k}(\mathbf{V}_k)$
12. $\alpha_k \leftarrow \eta \alpha_k$
13. **until** \mathbf{X}_{k+1} satisfies an acceptance criterion.
14. $k \leftarrow k + 1$
15. **until** some stopping criterion is satisfied.

A common acceptance criterion in line 13 requires the objective function to decrease; see [27] for details.

2.3. Debiasing

As is well known, the solutions obtained under 2OSCAR (and many other types of regularizers) are attenuated/biased in magnitude. Thus, it is common practice to apply *debiasing* as a postprocessing step; *i.e.*, the solutions obtained by, say the SpaRSA algorithm provides the structure/support of the estimate and the debiasing step recovers the magnitudes of the solutions. The debiasing method used in SpaRSA [27] is also adopted in this paper. Specifically, the debiasing phase solves

$$\begin{aligned} \hat{\mathbf{X}}_{\text{debias}} &= \arg \min_{\mathbf{X}} (\|\mathbf{A}\mathbf{X} - \mathbf{Y}\|_F^2) \\ \text{s.t. } \text{supp}(\mathbf{X}) &= \text{supp}(\tilde{\mathbf{X}}) \end{aligned} \quad (10)$$

where $\tilde{\mathbf{X}}$ is the estimate produced by the SpaRSA (or any other) algorithm and $\text{supp}(\mathbf{X})$ denotes the set of indices of the

non-zero elements of \mathbf{X} . This problem is solved by conjugate gradient procedure; see [27] for more details.

3. EXPERIMENTS

All the experiments were performed using MATLAB on a 64-bit Windows 7 PC with an Intel Core i7 3.07 GHz processor and 6.0GB of RAM. The performance of the different algorithms is assessed via the following five metrics, where \mathbf{E} is an estimate of \mathbf{X}):

- Mean absolute error, $\text{MAE} = \|\mathbf{X} - \mathbf{E}\|_1 / (nd)$;
- Mean square error, $\text{MSE} = \|\mathbf{X} - \mathbf{E}\|_F^2 / (nd)$;
- Position error rate,

$$\text{PER} = \sum_{i=1}^n \sum_{j=1}^d \left| |\text{sign}(\mathbf{X}_{(i,j)})| - |\text{sign}(\mathbf{E}_{(i,j)})| \right| / (nd).$$

- Elapsed time (TIME).

We consider the experiments on recovery of a 100×10 matrix \mathbf{X} with different styles of groups – blocks, lines and curved groups, consisting of positive and negative elements. The observed matrix \mathbf{Y} is generated by (1), in which the variance of the noise \mathbf{W} is $\sigma^2 = 0.16$. The sensing matrix \mathbf{A} is a 65×100 matrix with components sampled from the standard normal distribution. There are 100 nonzeros in the original 100×10 matrix, with values arbitrarily chosen from the set $\{-7, -8, -9, 7, 8, 9\}$ (Fig. 3).

Table 1. Results of metrics

Metrics	TIME (sec.)		MAE		MSE		PER
	yes	no	yes	no	yes	no	
FISTA	4.37	4.26	0.0784	0.477	2.45	0.202	0.1%
TwIST	5.10	4.45	0.0799	0.480	2.47	0.202	0.2%
SpaRSA	2.25	2.26	0.0784	0.477	2.44	0.202	0.0%
ADMM	6.65	6.60	0.0786	0.477	2.44	0.206	0.2%
SBM	6.32	6.22	0.0784	0.477	2.45	0.202	0.1%
PADMM	6.01	5.97	0.0762	0.456	2.42	0.182	0.0%

We run algorithms mentioned above (FISTA, TwIST, SpaRSA, SBM, ADMM, PADMM), with and without debiasing. The stopping condition is $\|\mathbf{X}_{k+1} - \mathbf{X}_k\| / \|\mathbf{X}_{k+1}\| \leq 0.001$, where \mathbf{X}_k represents the estimate at the k -th iteration. We set $\lambda_1 = 0.5$ and $\lambda_2 = 0.0024$. Other parameters are hand-tuned in each case for the best improvement in MAE. The recovered matrices are shown in Fig. 3 and the quantitative results are reported in Table 1.

We can conclude from Fig. 3 and Table 1 that the 2OSCAR criterion solved by proximal splitting algorithms with debiasing is able to accurately recover group-sparse matrices. Among the algorithms, the SpaRSA is the fastest, while the PADMM obtains the most accurate solutions.




Fig. 3. Original and recovered matrices

4. CONCLUSIONS

We have used the OSCAR regularizer to recover group-sparse matrix with arbitrary groups from compressive measurements. A matrix version of the OSCAR (2OSCAR) has been proposed and solved by six state-of-the-art proximal splitting algorithms: FISTA, TwIST, SpaRSA, SBM, ADMM and PADMM. Experiments on group-sparse matrix recovery show that the 2OSCAR regularizer solved by the SpaRSA algorithm has the fastest convergence, while the PADMM yields the most accurate estimates.

5. REFERENCES

- [1] S.F. Cotter, B.D. Rao, K. Engan, and K. Kreutz-Delgado, “Sparse solutions to linear inverse problems with multiple measurement vectors,” *IEEE Trans. on Signal Proc.*, vol. 53, pp. 2477–2488, 2005.
- [2] BD Rao and K. Kreutz-Delgado, “Sparse solutions to linear inverse problems with multiple measurement vectors,” in *8th IEEE Digital Signal Proc. Workshop*, 1998.
- [3] Z. Zhang and B.D. Rao, “Sparse signal recovery with temporally correlated source vectors using sparse bayesian learning,” *IEEE Jour. of Selected Topics in Signal Proc.*, vol. 5, pp. 912–926, 2011.
- [4] M. Yuan and Y. Lin, “Model selection and estimation in regression with grouped variables,” *Jour. of the Royal Statistical Society (B)*, vol. 68, pp. 49–67, 2005.
- [5] Z. Qin and D. Goldfarb, “Structured sparsity via alternating direction methods,” *The Jour. of Machine Learning Research*, vol. 98888, pp. 1435–1468, 2012.

- [6] L. Yuan, J. Liu and J. Ye, "Efficient methods for overlapping group lasso," *IEEE Trans. Patt. Analysis and Machine Intell.*, vol. 35, pp. 2104–2116, 2013.
- [7] J. Zhou, J. Chen, and J. Ye, "Clustered multi-task learning via alternating structure optimization," *Advances in Neural Inform. Proc. Systems*, vol. 25, 2011.
- [8] J. Zhou, L. Yuan, J. Liu, and J. Ye, "A multi-task learning formulation for predicting disease progression," in *17th ACM SIGKDD international Conf. on Knowledge Discovery and Data Mining*, 2011, pp. 814–822.
- [9] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski, "Structured sparsity through convex optimization," *Statistical Science*, vol. 27, pp. 450–468, 2012.
- [10] Y.C. Eldar and H. Bolcskei, "Block-sparsity: Coherence and efficient recovery," in *IEEE Int. Conf. on Acoustics, Speech and Signal Proc.*, 2009, pp. 2885–2888.
- [11] J. Huang, T. Zhang, and D. Metaxas, "Learning with structured sparsity," *The Jour. of Machine Learning Research*, vol. 12, pp. 3371–3412, 2011.
- [12] C.A. Micchelli, J.M. Morales, and M. Pontil, "Regularizers for structured sparsity," *Advances in Computational Math.*, pp. 1–35, 2010.
- [13] J. Mairal, R. Jenatton, G. Obozinski, and F. Bach, "Convex and network flow algorithms for structured sparsity," *The Jour. of Machine Learning Research*, vol. 12, pp. 2681–2720, 2011.
- [14] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Jour. of the Royal Statistical Society (B)*, pp. 267–288, 1996.
- [15] N. Simon, J. Friedman, T. Hastie, and R. Tibshirani, "The sparse-group lasso," *Jour. of Comput. and Graphical Statistics*, 2012, to appear.
- [16] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *Jour. of the Royal Statistical Society (B)*, vol. 67, pp. 301–320, 2005.
- [17] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight, "Sparsity and smoothness via the fused lasso," *Jour. of the Royal Statistical Society (B)*, vol. 67, pp. 91–108, 2004.
- [18] H.D. Bondell and B.J. Reich, "Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with OSCAR," *Biometrics*, vol. 64, pp. 115–123, 2007.
- [19] Z.J. Daye and X.J. Jeng, "Shrinkage and model selection with correlated variables via weighted fusion," *Computational Statistics & Data Analysis*, vol. 53, pp. 1284–1298, 2009.
- [20] S. Kim, K.A. Sohn, and E.P. Xing, "A multivariate regression approach to association analysis of a quantitative trait network," *Bioinformatics*, vol. 25, pp. i204–i212, 2009.
- [21] Q. Li and N. Lin, "The bayesian elastic net," *Bayesian Analysis*, vol. 5, pp. 151–170, 2010.
- [22] X. Shen and H.C. Huang, "Grouping pursuit through a regularization solution surface," *Jour. of the American Statistical Assoc.*, vol. 105, pp. 727–739, 2010.
- [23] L.W. Zhong and J.T. Kwok, "Efficient sparse modeling with automatic feature grouping," *IEEE Trans. on Neural Networks and Learning Systems*, vol. 23, pp. 1436–1447, 2012.
- [24] X. Zeng and M.A.T. Figueiredo, "Solving OSCAR regularization problems by proximal splitting algorithms," *arXiv preprint arxiv.org/abs/1309.6301*, 2013.
- [25] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Jour. on Imaging Sciences*, vol. 2, pp. 183–202, 2009.
- [26] J.M. Bioucas-Dias and M.A.T. Figueiredo, "A new twist: two-step iterative shrinkage/thresholding algorithms for image restoration," *IEEE Trans. on Image Proc.*, vol. 16, pp. 2992–3004, 2007.
- [27] S.J. Wright, R.D. Nowak, and M.A.T. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. on Signal Proc.*, vol. 57, pp. 2479–2493, 2009.
- [28] 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, pp. 1–122, 2011.
- [29] T. Goldstein and S. Osher, "The split Bregman method for l_1 -regularized problems," *SIAM Jour. on Imaging Sciences*, vol. 2, pp. 323–343, 2009.
- [30] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *Jour. of Math. Imaging and Vision*, vol. 40, pp. 120–145, 2011.
- [31] H.H. Bauschke and P.L. Combettes, *Convex analysis and monotone operator theory in Hilbert spaces*, Springer, 2011.
- [32] J. Barzilai and J.M. Borwein, "Two-point step size gradient methods," *IMA Jour. of Numerical Analysis*, vol. 8, pp. 141–148, 1988.