

**RECEIVED**
19 September 2024**ACCEPTED FOR PUBLICATION**
18 October 2024**PUBLISHED**
4 November 2024**Julianne Chung** and **Malena Sabaté Landman** *

Department of Mathematics, Emory University, Atlanta, 30322, GA, United States of America

* Author to whom any correspondence should be addressed.

E-mail: msabat3@emory.edu**Keywords:** flexible Krylov methods, group sparsity, image reconstruction, inverse problems

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Abstract

This paper introduces new solvers for efficiently computing solutions to large-scale inverse problems with group sparsity regularization, including both non-overlapping and overlapping groups. Group sparsity regularization refers to a type of structured sparsity regularization, where the goal is to impose additional structure in the regularization process by assigning variables to predefined groups that may represent graph or network structures. Special cases of group sparsity regularization include ℓ_1 and isotropic total variation regularization. In this work, we develop hybrid projection methods based on flexible Krylov subspaces, where we first recast the group sparsity regularization term as a sequence of 2-norm penalization terms using adaptive regularization matrices in an iterative reweighted norm fashion. Then we exploit flexible preconditioning techniques to efficiently incorporate the weight updates. The main advantages of these methods are that they are computationally efficient (leveraging the advantages of flexible methods), they are general (and therefore very easily adaptable to new regularization term choices), and they are able to select the regularization parameters automatically and adaptively (exploiting the advantages of hybrid methods). Extensions to multiple regularization terms and solution decomposition frameworks (e.g. for anomaly detection) are described, and a variety of numerical examples demonstrate both the efficiency and accuracy of the proposed approaches compared to existing solvers.

1. Introduction

Large-scale linear ill-posed inverse problems of the form

$$\mathbf{A}\mathbf{x}_{\text{true}} + \mathbf{e} = \mathbf{b}, \quad (1)$$

arise in the discretization of problems coming from various scientific and engineering applications, such as biomedical, atmospheric and medical imaging [1, 2]. In this form, \mathbf{x}_{true} is the unknown solution, \mathbf{b} is the measured data affected by unknown Gaussian white noise \mathbf{e} , and $\mathbf{A} \in \mathbb{R}^{m \times n}$ models the forward model. Given \mathbf{b} and \mathbf{A} , the goal of the inverse problem is to approximate \mathbf{x}_{true} , but computing solutions can be challenging especially for large-scale and ill-posed problems. More specifically, we are interested in cases where \mathbf{A} is ill-conditioned and has ill-determined rank, meaning that the singular values of \mathbf{A} decay and cluster at zero without a clear gap to indicate numerical rank. Due to the presence of noise in the measured data, the naive solution $\mathbf{A}^\dagger \mathbf{b}$ (where \mathbf{A}^\dagger is the Moore-Penrose pseudoinverse of \mathbf{A}) can differ significantly from the desired true solution $\mathbf{A}^\dagger(\mathbf{b} - \mathbf{e})$. This is due to noise amplification, as described in [1]. To obtain a meaningful approximation of the true solution, regularization is required. There are many forms of regularization, but the general idea of regularization is to include prior knowledge about the solution (e.g. smoothness assumptions, solution structure, or hard constraints) in the solution process.

In this work, we are interested in group sparsity regularization, which is a type of structured sparsity regularization, that promotes sparsity among pre-defined groups. Consider the general optimization framework where the goal is to solve the following variational regularization problem,

$$\min_{\mathbf{x}} \{ \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\Psi\mathbf{x}\|_{2,1} \}, \quad (2)$$

where $\lambda > 0$ is a regularization parameter, Ψ is an invertible matrix, and the $\ell_{2,1}$ norm for a vector \mathbf{z} whose elements belong to s different groups is defined as

$$\|\mathbf{z}\|_{2,1} = \sum_{i=1}^s \|\mathbf{z}_{g_i}\|_2, \quad (3)$$

with the vector \mathbf{z}_{g_i} denoting all components in \mathbf{z} belonging to the i -th group. The choice of groups promotes structural patterns in the unknown parameters, and the methodology presented in this paper can handle both non-overlapping and overlapping groups.

Group sparsity regularization has gained significant interest in the literature, due to recent applications to machine learning (e.g. in training of deep neural networks [3] and spatio-temporal problems [4]) and new theories (e.g. for strong group sparsity [5]). The benefits of using group sparsity regularization have been studied and demonstrated on various problems; however, various computational challenges remain that have hindered the use and extension of these group sparsity regularizers for large-scale problems. First, selecting an appropriate choice of the regularization parameter λ can be a difficult task, and most solution approaches require time-consuming tuning of this parameter. Second, for problems where the forward model matrix \mathbf{A} lacks an exploitable structure and may not even be explicitly constructed and stored, the only viable approach to solving problem (1) is to employ iterative methods. However, it is well-known that iterative methods, when applied to un-regularized ill-posed inverse problems, exhibit semiconvergence behavior, where solutions at later iterations become dominated by errors, so a good stopping criteria is critical [1, 6]. Third, in many scenarios, group sparsity on its own is not enough, and further refinement is needed. For example, regularizers that combine a group sparsity prior as well as an ℓ_1 (LASSO) regularizer [3, 7] result in problems of the form,

$$\min_{\mathbf{x}} \{ \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{x}\|_1 + \alpha \|\mathbf{x}\|_{2,1} \}, \quad (4)$$

for some $\alpha > 0$. One interpretation is that in addition to the nonzero components being clustered into groups, the nonzero groups themselves may be sparse.

Anomaly detection presents another scenario that requires going beyond the standard group sparsity problem (2). In this case, the solution can be represented as $\mathbf{x} = \xi + \mathbf{s}$, where ξ captures background smoothness and \mathbf{s} represents anomalous events [8, 9]. Such problems arise in the context of atmospheric inverse modeling, and these are often spatio-temporal problems where anomalies are consistent over time but sparse in terms of spatial location. Thus, group sparsity regularization can be used for \mathbf{s} with groups defined via time, while a spatio-temporal Gaussian prior can be used for ξ , e.g. $\xi \sim \mathcal{N}(0, \alpha^{-1}\mathbf{Q})$, where \mathbf{Q} is a symmetric positive definite (SPD) matrix. Assuming that \mathbf{R} is an SPD matrix representing the noise covariance matrix, we are interested in anomaly detection problems of the form,

$$\min_{\xi, \mathbf{s}} \{ \|\mathbf{A}(\xi + \mathbf{s}) - \mathbf{b}\|_{\mathbf{R}^{-1}}^2 + \alpha \|\xi\|_{\mathbf{Q}^{-1}}^2 + \lambda \|\mathbf{s}\|_{2,1} \}, \quad (5)$$

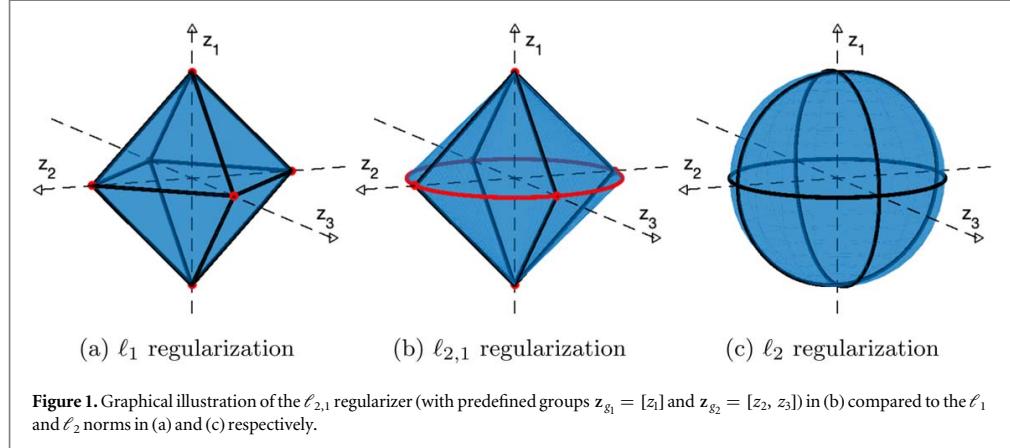
where $\|\mathbf{x}\|_{\mathbf{M}}^2 = \mathbf{x}^T \mathbf{M} \mathbf{x}$ for a given SPD matrix \mathbf{M} .

For both (4) and (5), a further challenge is to estimate the additional regularization parameter α . We will show in section 4.3 that the flexible Krylov methods for group sparsity can be extended to solve anomaly detection problems such as (5).

Contributions and overview

In this paper, we develop flexible Krylov methods, including hybrid and iteratively reweighted variants, for efficiently computing solutions to large-sale inverse problems with group sparsity regularization. Although Krylov methods have been considered for sparsity-promoting regularization terms [10–13], this work is the first to use flexible methods to address group sparsity regularizers. We address particular challenges that come with extensions to group sparsity regularizers, including handling overlapping groups, combining multiple regularization terms, and extending solution decomposition frameworks for anomaly detection. The proposed methods have the advantage of being general (and therefore very easily adaptable to new regularization term choices), very efficient (leveraging the advantages of flexible methods), and allowing for automatic and adaptive regularization parameter choices (exploiting the advantages of hybrid methods). The performance of this framework is shown through a variety of numerical examples showcasing both the efficiency of the proposed framework compared to existing solvers and showing novel applications of group sparsity.

The paper is organized as follows. In section 2 we provide a some background on group sparsity. Then in section 3 we describe how iterative reweighted norm schemes can be used to handle a group sparsity regularization term and describe various flexible hybrid projection methods, with particular focus on extensions to multiple regularization terms and to the solution decomposed problem. In section 4 we provide various examples from image deblurring, dynamic image deblurring and anomaly detection in atmospheric inverse modeling. Conclusions are provided in section 5.



2. Background on group sparsity

The group sparsity regularized problem (2) corresponds to a very general framework for regularization. In fact, the choice of specific groups leads to standard regularization terms. For example, ℓ_1 regularization can be obtained by considering n groups, with each group containing a single element, and 2-norm regularization can be obtained by considering one group containing all parameters.

Moreover, isotropic total variation (TV) regularization in any dimension (e.g. in 2D, 3D, and considering a temporal dimension) can be obtained by taking Ψ to be a discrete derivative operator and grouping all directional derivatives. Notice that a similar case to the combined regularized problem (4) is the elastic net regularizer, which combines an ℓ_1 and an ℓ_2 regularizer [14]; hence, the methods proposed in this paper to solve (4) can be straightforwardly adapted to solve the elastic net problem.

A major distinction of the group sparsity regularizer is the ability to select groups, where the choice of groups promotes structural patterns in the unknown parameters. For many problems, there is a natural choice for the groups that comes from the structure of the problem. For example, in source localization and anomaly detection where spatial-temporal images are desired, a natural grouping arises where each pixel at every time point constitutes a group. This has been studied, e.g. for image-based biochemical assay [15], low-dimensional nonlinear signal modeling [16], and classification [17]. In other applications, groups are defined by networks (e.g. deep neural networks) or by wavelet hierarchies [18].

The groups can be non-overlapping (e.g. (6) with $\Psi = \mathbf{I}$ corresponds to group LASSO), where a partition of the unknown vector is made, and the regularizer promotes sparsity for the entire group. Or, the groups can be overlapping, where a component of the unknown vector can belong to more than one group simultaneously. If the unknown parameters themselves do not lend themselves to natural grouping, an invertible matrix Ψ in (2) can be used to arrange the groups. Notice that (2) can be expressed as

$$\min_{\mathbf{z}} \{ \|\mathbf{A}\Psi^{-1}\mathbf{z} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{z}\|_{2,1} \}, \quad \text{for } \mathbf{x} = \Psi^{-1}\mathbf{z}. \quad (6)$$

A natural question is to how to interpret the group sparsity regularizer. In figure 1, we provide a graphical representation of non-overlapping group sparsity for a schematic 3D example, compared to standard ℓ_1 and ℓ_2 norms. The edges of the solid shapes in figures 1(a), (b) and (c), represent the different vectors $\mathbf{z} = [z_1, z_2, z_3] \in \mathbb{R}^3$ with $\|\mathbf{z}\|_1 = 1$, $\|\mathbf{z}\|_{2,1} = 1$ and $\|\mathbf{z}\|_2 = 1$, respectively, where the groups for the $\ell_{2,1}$ norm have been taken to be $\mathbf{z}_{g_1} = [z_1]$ and $\mathbf{z}_{g_2} = [z_2, z_3]$. The geometry of the norm ball is directly related to the solution of (6).

In particular, the positions of the singularities are different for each norm. For example, the ℓ_1 norm promotes sparsity in the individual coefficients. To see this, imagine a plane corresponding to the set of solutions with equal discrepancy (i.e. giving the same values for the fit-to-data term $\|\mathbf{A}\Psi^{-1}\mathbf{z} - \mathbf{b}\|_2^2$). The intersection of this plane with the octahedron in figure 1(a) will most likely occur at a corner (i.e. denoted by the red dots in figure 1(a)), which corresponds to a sparse solution. On the contrary, the differentiable-everywhere ℓ_2 -norm which represents a sphere in \mathbb{R}^3 (see figure 1(c)) does not favor any particular direction (i.e. does not promote sparsity in the solution). The $\ell_{2,1}$ norm is represented as a double-cone in figure 1(b), and it provides a natural balance between the ℓ_1 and ℓ_2 norms. Since the two cones are attached at their bases, $\ell_{2,1}$ regularization promotes sparsity, but only at the group level, i.e. either z_1 or both z_2 and z_3 . Notice the absence of singularities at

standard unit vectors \mathbf{e}_2 and \mathbf{e}_3 , such that the likely points of intersection (e.g. with a plane) would occur at the points denoted in red.

For overlapping group sparsity, a graphical illustration is not as intuitive and depends on the choice of overlapping weights. Nevertheless, overlapping group sparsity can be very useful in some applications (e.g. tree structured groups, contiguous groups, and directed-acyclic-graph groups). It has been shown in [19] that structured sparsity assumptions can improve computational efficiency and accuracy in the solution in an analysis analogous to traditional compressive sensing for sparsity. This can be exploited, for example, when using tree-based image representations, see e.g. [20] for a characterization using the block-DCT transform associated with JPEG, or [18] for image representations using wavelets. In this setting, sub-trees are formed by grouping hierarchically different scale representations at the same location, so for natural images it is reasonable to assume that the image would be zero at all resolutions in a given point; therefore, one can group parents and children together in different ways. For example, this is exploited in [21], where an algorithm based on preconditioned conjugate gradient with restarts is presented, and [22], where group-wise soft thresholding is used to deal with the $\ell_{2,1}$ regularization term.

Many of the challenges to computing solutions to (2) mimic those in LASSO and compressive sensing (e.g. a discontinuous first derivative at zero from the 1-norm, causing challenges for optimization algorithms). Various methods have been proposed to solve such problems, ranging from convex relaxations (e.g. iteratively reweighted schemes) and nonconvex optimization methods (e.g. ADMM). A main caveat of all of these methods is the need to select the regularization parameter *a priori*. More recently, flexible Krylov methods have been shown to be powerful alternatives to solve ℓ_1 regularization problems (see [11, 23]) compared to other classical solvers such as suitable versions of FISTA [24] or SpaRSA [25], or using iteratively reweighted norms (maybe) in combination with Krylov methods (see, e.g. [26, 27]). Similarly, flexible Krylov methods have been used for TV regularization (see [13, 28]) and proved competitive against the most popular solvers for this problem such as proximal gradient methods [29], primal-dual methods [30], split Bregman methods [31] or methods based on iteratively reweighted norms [32]. Another example of modified TV using group sparsity can be found in [4], where group sparsity is used to promote piece-wise constant structures in space and time, sparsity in gradient images shared across time instances, but not among the group. This is solved using a majorization-minimization approach based on a generalized Krylov subspace approach [33, 34]. For non-overlapping groups, accelerated proximal gradient methods have been considered with added non-negativity constraints to solve (2) in [15] and alternating direction method of multipliers was used in [7].

It is beyond the scope of this paper to handle cases where the groups are not known *a priori*. However, the methodology presented in section 3 could be used in combination with a strategy to determine (or modify) groups on-the-fly once a new approximation of the solution is available, since the weights are iteration dependent. Similarly, the flexible Krylov methods can be easily generalised to $\ell_{2,p}$ by using suitably defined weights. Moreover, if memory requirements are a problem, iteratively reweighted schemes with fixed preconditioning can be used in an inner-outer scheme.

3. Flexible Krylov methods for group sparsity

In this section, we describe flexible Krylov methods and hybrid variants for solving the group sparsity regularized problem (2) for both overlapping and non-overlapping groups, and we develop extensions to solve (4) and (5). We begin in section 3.1 with a description of the iteratively reweighted norm approach.

Then in section 3.2, we describe flexible Krylov methods for efficiently generating a single solution basis for approximately solving the re-weighted least-squares subproblems. A method for automatically selecting the regularization parameter within the hybrid framework is described.

3.1. Iteratively reweighted norm schemes

Traditional iteratively reweighted (IRW) schemes are majorization-minimization methods that involve constructing a sequence of least-squares problems approximating a variational regularization problem of interest, which are then solved at each iteration. These methods intrinsically rely on the interpretation of a given ℓ_p norm as a non-linear weighted ℓ_2 norm (for $0 < p \leq 2$), such that evaluating the weights at an available approximation of the solution \mathbf{x}_{k-1} can be used to obtain a quadratic tangent majorant of the original problem at \mathbf{x}_{k-1} . Recall that a quadratic tangent majorant of a non-quadratic function at a point $\bar{\mathbf{z}}$ is a quadratic upper bound of the original function constructed such that both the value of the function and the value of the gradient of both functions coincide at $\bar{\mathbf{z}}$. More specifically, for $\Phi(\mathbf{z}) = \frac{1}{p} \|\mathbf{z}\|_p^p = \frac{1}{p} \|\mathbf{W}(\mathbf{z})\mathbf{z}\|_2^2$, then $\Phi^{\bar{\mathbf{z}}}(\mathbf{z}) = \frac{1}{2} \|\mathbf{W}(\bar{\mathbf{z}})\mathbf{z}\|_2^2 + (\frac{1}{p} - \frac{1}{2}) \|\mathbf{W}(\bar{\mathbf{z}})\bar{\mathbf{z}}\|_2^2$ is a quadratic tangent majorant of $\Phi(\mathbf{z})$ at $\bar{\mathbf{z}}$. This can be explicitly checked as

$\Phi^{\bar{z}}$ is quadratic, $\Phi^{\bar{z}}(\bar{z}) = \Phi(\bar{z})$, the gradients are equal at \bar{z} , i.e. $\nabla\Phi^{\bar{z}}(\bar{z}) = \nabla\Phi(\bar{z})$, and $\Phi^{\bar{z}}(z)$ is an upper bound of $\Phi(z)$ for all z (see, e.g. [35]).

IRW norm schemes can be used to solve minimization problems involving $\ell_{2,1}$ norm terms. Assume the components of $z \in \mathbb{R}^n$ belong to s different (possibly overlapping) groups and define z_{g_i} to be the vector constructed by concatenating the components in z belonging to the i -th group. Let G_i be the indexes of the (possibly multiple) groups where the component z_j of z belongs to and, with a small notation abuse, let $z_j \in z_{g_i}$ be the components of z that belong to the i -th group. Then, we can define a diagonal weighting matrix (as a function of z) as

$$W_{j,j}(z) = \sqrt{\sum_{i \in G_j} \frac{1}{\|z_{g_i}\|_2^2}}, \quad (7)$$

such that the $\ell_{2,1}$ norm defined in (3) can be written as a reweighted ℓ_2 norm:

$$\begin{aligned} \|W(z)z\|_2^2 &= \sum_{j=1}^n \left(\sum_{i \in G_j} \frac{1}{\|z_{g_i}\|_2^2} z_j^2 \right) = \sum_{i=1}^s \left(\sum_{z_j \in z_{g_i}} \frac{1}{\|z_{g_i}\|_2^2} z_j^2 \right) \\ &= \sum_{i=1}^s \frac{\|z_{g_i}\|_2^2}{\|z_{g_i}\|_2^2} = \sum_{i=1}^s \|z_{g_i}\|_2 = \|z\|_{2,1}. \end{aligned} \quad (8)$$

The solution-dependent non-linear weights defined in (8), can be then used to re-write the original regularization problem (6) with group sparsity regularization as

$$\min_z \{ \|A\Psi^{-1}z - b\|_2^2 + \lambda \|W(z)z\|_2^2 \} \quad \text{for } x = \Psi^{-1}z. \quad (9)$$

Using an IRW scheme corresponds to evaluating (and fixing) the weights (7) at each (outer) iteration using the available approximation of the solution coming from the previous problem in the sequence. However, this can be numerically unstable due to divisions by zero when one has to evaluate the weights on a vector with groups whose components are all 0 valued. This situation is in fact expected, as we are assuming sparsity in the groups, and is caused by the lack of smoothness of the functional in (6) at such vectors with 0 valued components. Thus, we consider a smooth approximation of the original problem, and the following weights are considered instead of the ones defined in (7):

$$\widetilde{W}_{j,j}(z) = \sqrt{\sum_{i \in G_j} \frac{1}{\|z_{g_i}\|_2^2 + \tau^2}}, \quad (10)$$

so that $W_k = \widetilde{W}(z_{k-1}) \approx W(z_{k-1})$. This yields the following sequence of least-squares problems approximating (6):

$$\min_z \{ \|A\Psi^{-1}z - b\|_2^2 + \lambda \|W_k z\|_2^2 \} \quad \text{for } x = \Psi^{-1}z. \quad (11)$$

Note that this minimization is equivalent to minimizing a sequence of quadratic tangent majorants of a smoothed version of (6), where multiplicative constants have been absorbed by λ and additive constants have been disregarded. Moreover, note that one could transform problem (11) to the following equivalent expression using a solution-dependent right preconditioner (or, equivalently, a suitable change of variables):

$$\min_s \{ \|A\Psi^{-1}W_k^{-1}s - b\|_2^2 + \lambda \|s\|_2^2 \} \quad \text{for } x = \Psi^{-1}z = \Psi^{-1}W_k^{-1}s. \quad (12)$$

Expression (12) motivates the use of flexible Krylov methods which allow for iteration-dependent preconditioning, see e.g. [11, 23] and references therein, and will be explained in detail the following section. Note that reweighting schemes can also be used to solve (5) using (8) and (10) to derive analogous expressions to (9), (11) and (12), see [9].

Moreover, we can handle the minimization (4) with multiple regularization terms by defining two solution-dependent diagonal weight matrices W_1 and W_2 such that (4) can be written as the following reweighted problem:

$$\min_x \{ \|Ax - b\|_2^2 + \lambda \|W_1(x)x\|_2^2 + \alpha \|W_2(x)x\|_2^2 \} \quad (13)$$

or, equivalently,

$$\min_x \left\{ \|Ax - b\|_2^2 + \left\| \begin{bmatrix} \sqrt{\lambda}W_1(x) \\ \sqrt{\alpha}W_2(x) \end{bmatrix} x \right\|_2^2 \right\}. \quad (14)$$

For multiple regularization terms, a point of concern especially in the context of iterative projection methods, is the estimation of regularization parameters. An alternative is to define a functional representation

for the relationship between regularization parameters [36, 37]. For example, define $\alpha = \tau_\lambda^2 \lambda$ for some constant τ_λ . Then, let

$$\mathbf{L} = \begin{bmatrix} \sqrt{\lambda} \mathbf{W}_1(\mathbf{x}) \\ \sqrt{\alpha} \mathbf{W}_2(\mathbf{x}) \end{bmatrix} = \sqrt{\lambda} \begin{bmatrix} \mathbf{W}_1(\mathbf{x}) \\ \tau_\lambda \mathbf{W}_2(\mathbf{x}) \end{bmatrix} = \sqrt{\lambda} \mathbf{Q}^L \mathbf{D} \quad (15)$$

be the QR factorization of \mathbf{L} . Notice that since \mathbf{L} is a concatenation of diagonal matrices which are invertible (by definition), \mathbf{L} is guaranteed to have orthogonal columns and is full column rank. Thus, the diagonal elements of \mathbf{D} can be easily obtained by taking the norms of the columns of \mathbf{L} . Next, notice that \mathbf{D} is a diagonal matrix that depends on \mathbf{x} and \mathbf{Q}^L contains orthonormal columns, so we get the IRW regularized problem,

$$\min_{\mathbf{x}} \{ \| \mathbf{A} \mathbf{x} - \mathbf{b} \|_2^2 + \lambda \| \mathbf{D}(\mathbf{x}) \mathbf{x} \|_2^2 \}. \quad (16)$$

Flexible hybrid methods can be used to solve (16), and extensions to more than two regularization terms is straightforward. Numerical results demonstrating this approach are provided in section 4.2.

3.2. Flexible Krylov methods

Flexible Krylov methods are a subclass of Krylov methods that allow for iteration-dependent right preconditioning. The methodology presented here can be used generally for any flexible Krylov method, and in particular we will present three different methods to showcase the use of these algorithms in different applications.

One interpretation of flexible Krylov methods where the preconditioning is motivated by a reweighted norm in the regularization term is that some information about the solution is embedded in the solution space [38]. Moreover, this is achieved efficiently by building a single solution space, thereby avoiding inner-outer schemes.

Since the projected space already contains information about the (variational) regularization term, one can also set $\lambda = 0$ in (12) and obtain a suitably regularized solution equipped with early stopping. Without appropriate stopping, however, the solution will approach the naive solution $\mathbf{A}^\dagger \mathbf{b}$. This is the original approach followed in flexible Krylov methods and will be denoted with the prefix (F). Another approach is to add regularization in the projected problem in order to avoid semi-convergence. With a suitable regularization parameter, which can be adaptively set at each iteration, this method delivers a good reconstruction. However, note that the regularization parameter for the projected problems might not be suitable for the original problem (6). This can be understood since, in the limit, the solution obtained with this method will converge to the solution of the fit-to-data term in (6) with added Tikhonov regularization. This is usually noted with the prefix (hybrid-F) and will be the preferred method used throughout the paper. For a more detailed explanation, see [23]. Lastly, problem (6) can be projected onto the flexible Krylov subspace. Assuming no break-down has happened at $k = \min(m, n)$, one can extend the algorithm for $k > \min(m, n)$ by updating the weights in (6). In this case, the solution provided by a flexible method following this scheme will converge to the solution of the (smoothed version of) (6). This comes at the cost of a QR factorization of a tall and skinny matrix at each iteration and will be denoted with the prefix (IRW-F). The convergence proof for this scheme can be found in [23].

3.2.1. Hybrid-FGMRES

Given $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{b} \in \mathbb{R}^n$, iteration independent right preconditioner $\Psi^{-1} \in \mathbb{R}^{n \times n}$ and iteration dependent right preconditioning matrices $\mathbf{W}_k^{-1} \in \mathbb{R}^{n \times n}$, the flexible Arnoldi method is a process that, at the k th iteration, constructs vectors \mathbf{v}_{k+1} and $\mathbf{z}_k = \mathbf{W}_k^{-1} \mathbf{v}_k$ such that

$$\mathbf{A} \Psi^{-1} \mathbf{Z}_k = \mathbf{V}_{k+1} \mathbf{H}_k, \quad (17)$$

where $\mathbf{H}_k \in \mathbb{R}^{(k+1) \times k}$ is upper Hessenberg, $\mathbf{Z}_k = [\mathbf{z}_1 \dots \mathbf{z}_k] \in \mathbb{R}^{n \times k}$, and $\mathbf{V}_{k+1} = [\mathbf{v}_1 \dots \mathbf{v}_{k+1}] \in \mathbb{R}^{n \times (k+1)}$ has orthonormal columns. Here, we assume that $\mathbf{x}_0 = 0$, so $\mathbf{v}_1 = \mathbf{b} / \|\mathbf{b}\|_2$. A detailed algorithm for flexible Arnoldi can be found in e.g. [11].

Using hybrid flexible GMRES (hybrid-FGMRES) to promote group sparsity regularization involves projecting the least-squares problem $\| \mathbf{A} \mathbf{x} - \mathbf{b} \|_2^2$ onto a modified (flexible) Krylov subspace using the decomposition in (17) and adding regularization. Then, the solution at each iteration k can be computed as

$$\mathbf{x}_k = \Psi^{-1} \mathbf{z}_k = \Psi^{-1} \mathbf{Z}_k \mathbf{y}_k \quad \text{where} \quad \mathbf{y}_k = \arg \min_{\mathbf{y}} \{ \| \mathbf{H}_k \mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1 \|_2^2 + \lambda \| \mathbf{y} \|_2^2 \}, \quad (18)$$

where $\mathbf{e}_1 \in \mathbb{R}^{k+1}$ is the first column of the identity matrix of order $k+1$ and we have used the fact that $\mathbf{V}_{k+1} \|\mathbf{b}\|_2 \mathbf{e}_1 = \mathbf{b}$. Note that the iteration dependent preconditioners are built using $\mathbf{W}_k = \widetilde{\mathbf{W}}(\mathbf{Z}_{k-1} \mathbf{y}_{k-1}) = \widetilde{\mathbf{W}}(\mathbf{z}_{k-1})$ as defined in (10), which depend on the solution computed at the previous iteration.

3.2.2. Hybrid-FLSQR

For rectangular matrices $A \in \mathbb{R}^{m \times n}$, we can use a flexible LSQR approach, where for $b \in \mathbb{R}^m$, an iteration independent right preconditioner $\Psi^{-1} \in \mathbb{R}^{n \times n}$ and iteration dependent right preconditioning matrices $W_k^{-1} \in \mathbb{R}^{n \times n}$, the flexible Golub-Kahan decomposition computes at the k th iteration vectors u_{k+1}, v_k and $z_k = W_k^{-1}v_k$ such that

$$A\Psi^{-1}Z_k = U_{k+1}M_k \quad \text{and} \quad \Psi^{-\top}A^\top U_{k+1} = V_{k+1}S_{k+1}, \quad (19)$$

where, if $x_0 = 0$, $u_1 = b/\|b\|_2$. Here $M_k \in \mathbb{R}^{(k+1) \times k}$ is upper Hessenberg, $S_{k+1} \in \mathbb{R}^{(k+1) \times (k+1)}$ is upper triangular and both $V_{k+1} \in \mathbb{R}^{n \times (k+1)}$ and $U_{k+1} \in \mathbb{R}^{m \times (k+1)}$ contain orthonormal columns. For more details see, e.g. [39].

Using hybrid flexible LSQR (hybrid-FLSQR) to promote group sparsity regularization involves projecting the least-squares problem $\|Ax - b\|_2^2$ onto a modified (flexible) Krylov subspace using the decomposition in (19) and adding regularization to the projected problem. Then, we can compute an approximation of the solution at each iteration k as

$$x_k = \Psi^{-1}z_k = \Psi^{-1}Z_k y_k \quad \text{where} \quad y_k = \arg \min_y \{ \|M_k y - \|b\|_2 e_1\|_2^2 + \lambda \|y\|_2^2 \}. \quad (20)$$

In particular, we are interested in the case where $W_k = \widetilde{W}(Z_{k-1}y_{k-1}) = \widetilde{W}(z_{k-1})$ and therefore it depends on the solution computed at the previous iteration.

3.2.3. Hybrid-SD

With appropriately defined weights, both hybrid-GMRES and hybrid-FLSQR can be used to solve group sparsity regularized problems (2) and (4), but a different projection approach is needed to handle the solution decomposition problem (5) (e.g. for anomaly detection). In this work, we use the Flexible Generalized Golub-Kahan (FGGK) approach described in [9] to generate a basis for the solution, where the weights are determined from the group sparsity regularizer.

Given $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, SPD matrices $Q \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{m \times m}$, and iteration-dependent right preconditioning matrices $W_k^{-1} \in \mathbb{R}^{n \times n}$, the flexible, generalized Golub-Kahan iterative process generates vectors $v_k, z_k = W_k^{-1}v_k$, and u_{k+1} such that at iteration k ,

$$[AQ \quad A]\hat{Z}_k = U_{k+1}M_k \quad \text{and} \quad A^\top R^{-1}U_{k+1} = V_{k+1}S_{k+1}, \quad (21)$$

where

$$\hat{Z}_k = \begin{bmatrix} v_1 & \dots & v_k \\ z_1 & \dots & z_k \end{bmatrix} = \begin{bmatrix} V_k \\ Z_k \end{bmatrix}, \quad (22)$$

and $u_1 = b/\|b\|_R^{-1}$ and $v_1 = A^\top R^{-1}u_1$ if $x_0 = 0$. Note that $M_k \in \mathbb{R}^{(k+1) \times k}$ is upper Hessenberg, $S_{k+1} \in \mathbb{R}^{(k+1) \times (k+1)}$ is upper triangular, and in exact arithmetic, the columns of $U_{k+1} \in \mathbb{R}^{m \times (k+1)}$ and $V_{k+1} \in \mathbb{R}^{n \times (k+1)}$ satisfy the orthogonality conditions,

$$U_{k+1}^\top R^{-1}U_{k+1} = I_{k+1}, \quad V_{k+1}^\top QV_{k+1} = I_{k+1}. \quad (23)$$

In order to use the FGGK process to solve (5), we first consider the sequence of reweighted least-squares problems,

$$\min_{\psi, s} \{ \|AQ\psi + As - b\|_{R^{-1}}^2 + \alpha \|\psi\|_Q^2 + \lambda \|W(s)s\|_2^2 \}, \quad (24)$$

where $\xi = Q\psi$ and W is defined as in (7). We first project the objective function $\|AQ\psi + As - b\|_{R^{-1}}^2 + \alpha \|\psi\|_Q^2$ onto a modified (flexible) Krylov subspace using the decomposition in (21), then add regularization to the projected problem, such that an approximation of the solution can be found at each iteration k as:

$$x_k = \xi_k + s_k = QV_k y_k + Z_k y_k \quad \text{where} \quad (25)$$

$$y_k = \arg \min_y \{ \|M_k y - \|b\|_{R^{-1}} e_1\|_2^2 + \alpha \|y\|_2^2 + \lambda \|R_{W,k} y\|_2^2 \}, \quad (26)$$

where $Z_k = Q_{W,k}R_{W,k}$ is a thin QR factorization that can be computed via efficient updates [9]. Contrary to the additional regularization terms added to (18) and (20), a non-standard Tikhonov regularizer is included here to distinguish the contributions from the two solution components. Notice that a decomposition of the solution x_k is available, i.e. ξ_k is a reconstruction of the smooth component and s_k is a reconstruction of the group-sparse component (e.g. containing anomalies). Moreover, in this case, the weights depend only on a part of the approximate solution at each iteration k , i.e. $W_k = \widetilde{W}(s_{k-1}) = \widetilde{W}(Z_{k-1}y_{k-1})$. A detailed explanation of the algorithm developed to obtain (21), originally proposed under the name of solution decomposition hybrid projection approach (sdHybr), can be found in [9].

Table 1. Notation for the solvers: **projection - ALGORITHM - GROUPS.**

projection	ALGORITHM	GROUPS
none	FLSQR	none
hybrid	FGMRES	G
IRW	SD	G1, G2
	RCG	C

3.2.4. Regularization parameter choice criteria

One of the main advantages of flexible Krylov methods is the ability to incorporate regularization on the projected problem (IRW-F and hybrid-F), and hence to adaptively set the regularization parameter(s) λ_k at each iteration.

Note that this is a very efficient approach. Differently to other methods, where the use of regularization parameter choice criteria requires the (full) solve of several regularized problems, here the regularization parameters are chosen on-the-fly throughout the iterations, only using information available from the projected problem and without requiring any extra matrix-vector products with the system matrix. This is also the approach followed by the more general class of hybrid methods, see, e.g. the review [40].

In particular, we consider the use of the discrepancy principle (DP)[41] for the projected problem, i.e. finding the regularization parameters $\lambda \geq 0$ such that

$$\|Ax_k(\lambda) - b\|_2 = \|Az_k y_k(\lambda) - b\|_2 = \eta \|e\|_2. \quad (27)$$

Note that $\lambda = \lambda$ for the minimization in (18) and (20), but $\lambda = [\lambda, \alpha]$ for the optimization in (26). The approximate solution computed by any given algorithm at iteration k is given by $x_k(\lambda)$, with y_k being the corresponding coefficients in the projected space. Using the DP requires a good estimate of the noise level and has been considered within a multi-parameter regularization context, e.g. see [42–44]. Alternatively, other methods can be used seamlessly, see, e.g. [40, 45].

4. Numerical examples

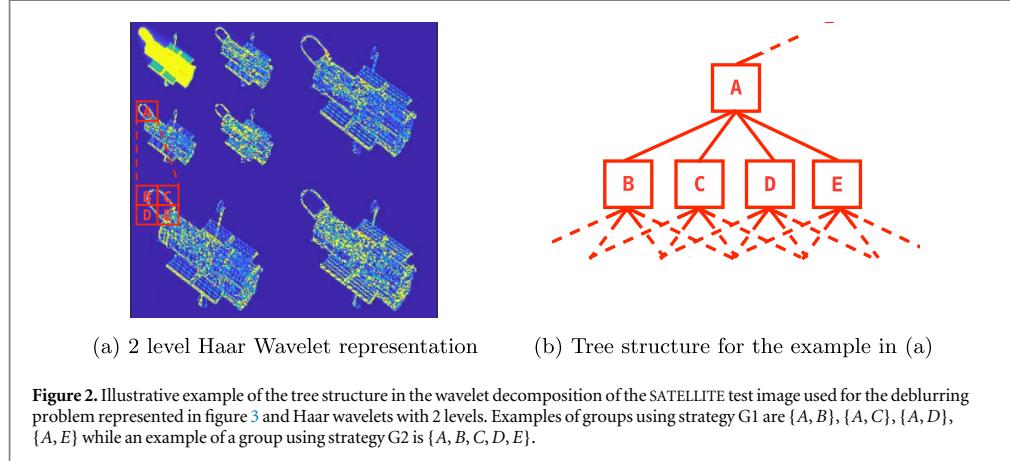
In this section three imaging examples are presented to demonstrate the performance of flexible Krylov methods for group sparsity regularization. In order to showcase the generality of the described approach, the flexible schemes are used for different grouping strategies including both overlapping and non-overlapping groups and in combination with different solvers.

The notation used to describe the solvers in this section can be found in table 1, where group sparsity regularization is indicated by appending a ‘G’ suffix at the end of the method. In the case where more than one grouping strategy has been compared, different grouping strategies are denoted ‘G1’, ‘G2’, etc. Moreover, if both ℓ_1 and $\ell_{2,1}$ are considered simultaneously, we use ‘C’ to denote the combined approach. If no group suffix is appended for a flexible method, then this corresponds to ℓ_1 regularization. Recall that the projection types are explained in section 3.2. And note that hybrid solution decompose (SD) [9] and restarted conjugate gradient (RCG) for group sparsity [21] are used as a comparison with previous works.

The first experiment exploits group sparsity patterns in wavelet coefficients by considering the natural tree structure of wavelet decompositions. Since non-leaf elements belong to more than one branch, this example showcases the use of the new scheme for regularization with overlapping groups. A deblurring problem is considered and solved using different versions of the new scheme involving FLSQR, paying particular attention to the different projection types and recalling the differences between them while showcasing their practical performance.

The second experiment corresponds to a dynamic deblurring problem. Since this is a spatio-temporal problem displaying a spatial sparsity pattern but smoothness in time, the groups are chosen to contain each pixel across all time points, so this is an example of non-overlapping group sparsity. In particular, we focus on the difference between enforcing ℓ_1 regularization and $\ell_{2,1}$ regularization using hybrid flexible methods, and we present results for the new method combining both regularization terms. Since this is a symmetric problem, the performance of both methods based on FLSQR and FGMRES is tested.

The third experiment concerns an anomaly detection problem in atmospheric inverse modeling. This corresponds to a realistic example, and our aim is to show the potential of the new schemes in real-world applications. Since this experiment also involves spatio-temporal images that are sparse in space but not time, the groups are built analogously to the second experiment. Additionally, since the solution in this experiment is



modelled as the sum of a smooth (background) component and a group-sparse (anomalies) component, hybrid-SD-G is used in this experiment.

For all the experiments, τ in (10) is taken to be 10^{-10} . However, empirical observations demonstrate that the algorithms are robust to this choice as long as τ is significantly smaller than the average pixel intensity and kept above machine precision. Moreover, the regularization parameter for all hybrid and iteratively-reweighted methods is chosen across the iterations using the DP defined in (27).

4.1. Deblurring example with wavelet sparsity patterns

Natural images have been shown to be well described by wavelet representations [46], Chapter 9. Moreover, wavelets have a natural tree-structure where each non-leaf wavelet coefficient for a given orientation (diagonal, vertical and horizontal) has four children coefficients corresponding to the same orientation at a finer scale. This is explained in, e.g. [18]. There are of course different groupings than one could develop. For example, in this paper, both 2-elements groups of each children with their parent (G1) or 5-elements groups of each parent with all their children (G2) are considered as suggested in [18]. In figure 2 an upscaled example for a wavelet decomposition with two levels is shown to illustrate a case of parent-child relationship in the wavelet domain. Note that figure 2 is intended as an illustration, but in reality the tree structure is considered at the pixel scale, i.e. where the red boxes in figure 2 correspond to individual pixels, to match the wavelet coefficients. Thus, at each level, each given pixel in the wavelet representation corresponding to a non-leaf element in the tree is the parent of 4 pixels in a finer level representing the same area of the original image (at that given orientation).

For this deblurring example, medium Gaussian blur is generated with IR Tools [47] with added Gaussian noise at noise level $\|\mathbf{e}\|/\|\mathbf{Ax}_{\text{true}}\| = 0.05$. Here the true solution $\mathbf{x}_{\text{true}} \in \mathbb{R}^{65,536}$ contains 256×256 pixels and corresponds to the SATELLITE test image in IR Tools and can be observed in figure 3 along with the corresponding blurred and noisy measurements. One can observe that the true image contains regions with details (and edges) at different scales, while other areas of the image contain many zeroes and hence do not have information at any scale. Therefore, group sparsity in the wavelet coefficients should provide suitable regularization.

Exploiting group sparsity in the wavelet coefficients corresponds to choosing Ψ in (2) to be an orthogonal wavelet transform and choosing the groups according to G1 or G2 as defined at the beginning of this subsection (and illustrated in figure 2). Note that both grouping schemes correspond to overlapping group sparsity regularization. The results for these experiments can be found in figure 4, where we provide relative reconstruction error norms, computed as $\|\mathbf{x}_k - \mathbf{x}_{\text{true}}\|_2 / \|\mathbf{x}_{\text{true}}\|_2$ where \mathbf{x}_k is the reconstruction at the k th iteration.

We provide a comparison against the algorithm presented in [21] (following their original grouping structure G1). Note that we have not included their proposed preconditioning, nor their starting guess ($\mathbf{x}_0 = 0$ is used instead of $\mathbf{A}^\top \mathbf{b}$), since those are highly problem-dependent choices and their approach is not suitable for this example. The regularization parameter for the hybrid flexible and IRW flexible methods is chosen at each iteration using the DP (27) with safety parameter $\eta = 1.01$. For the RCG algorithm [21], the regularization parameter is chosen to be the one computed with the hybrid flexible methods at the end of the iterations. It is worth mentioning that RCG is better suited to other forward models (such as MRI), but flexible Krylov methods have a comparable performance and allow the (semi-automatic) computation of the regularization parameter on-the-fly.

One can observe that flexible Krylov methods display a fast convergence in comparison to RCG. Different grouping strategies, displayed in figures 4(a) and (b) respectively, have similar performances (note that both

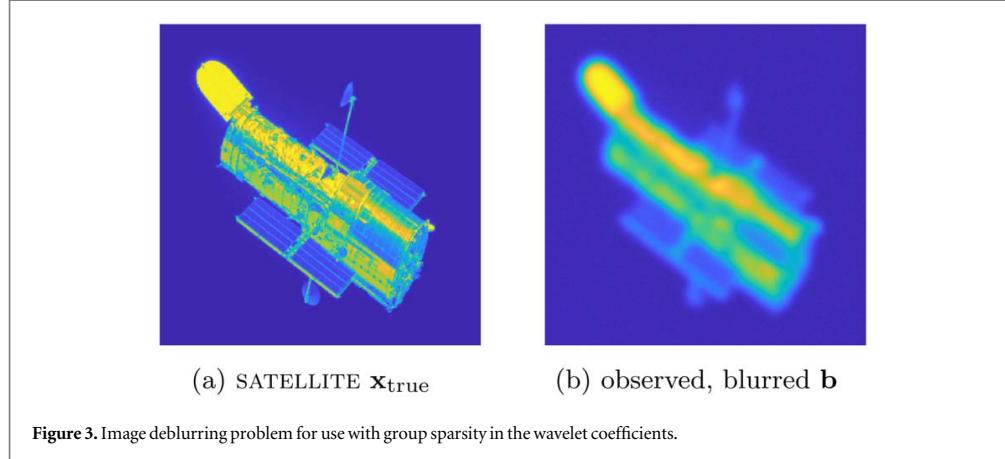


Figure 3. Image deblurring problem for use with group sparsity in the wavelet coefficients.

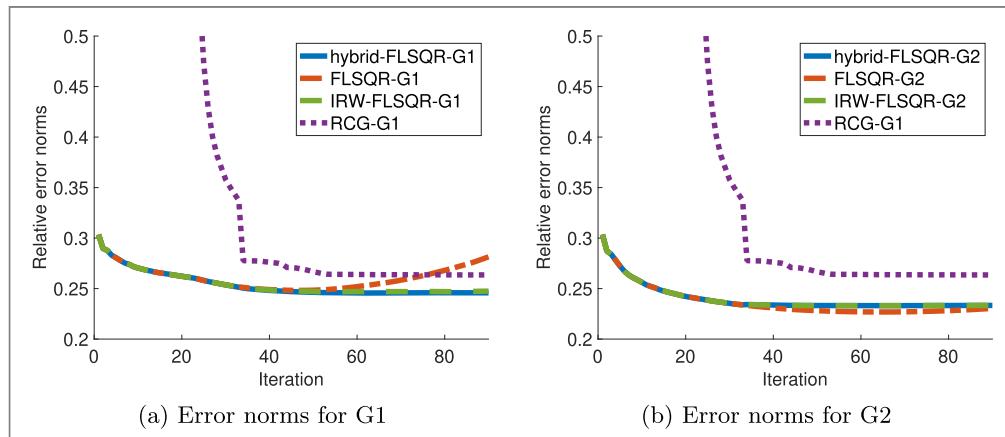
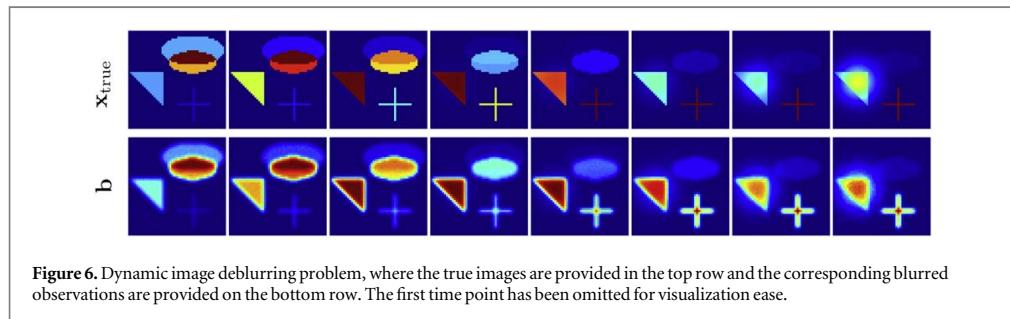
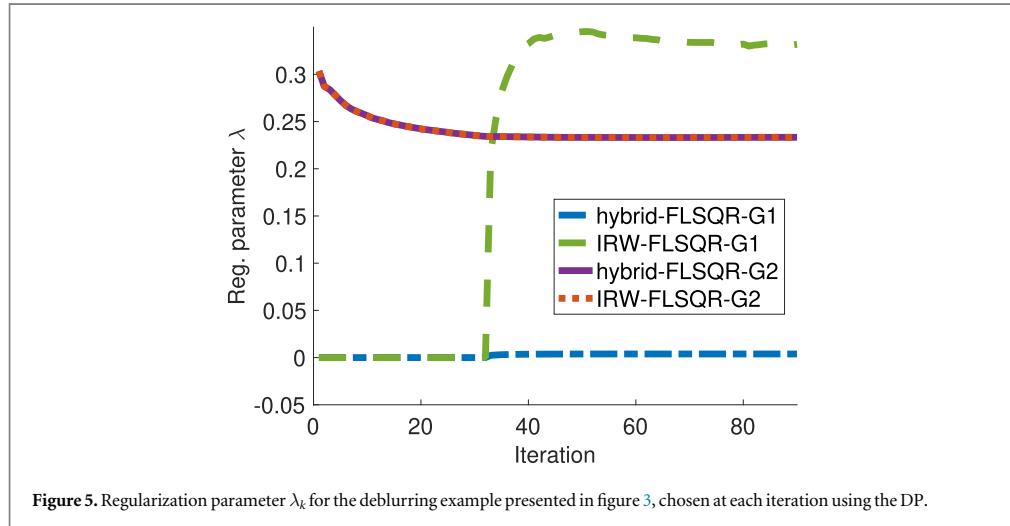


Figure 4. Relative reconstruction error norms for the methods based on FLSQR for the deblurring example presented in figure 3 with wavelet group sparsity. The regularization parameter has been chosen at each iteration using the DP.

error norm plots have the same scale for the axis). Moreover, one can observe that flexible Krylov methods without regularization in the projected problem (FLSQR-G1 and FLSQR-G2) display semi-convergence, and that the severity of it depends on the grouping strategy. To be able to use these algorithms in practice, one should use appropriate early stopping. On the contrary, methods with explicit regularization in the projected problem (hybrid-FLSQR-G and IRW-FLSQR-G for both grouping strategies), using the DP as a parameter choice criterion at each iteration, display a stabilization of the error norm throughout the iterations. It is worth noting that the regularization parameter for hybrid and IRW methods does not necessarily need to be the same as that for the original problem, as explained in section 3.2. In fact, for the grouping strategy G1, very similar error norms are achieved for the two methods using very different regularization parameters. This can be observed in figure 5. Recall that hybrid methods are cheaper per iteration than iteratively-reweighted methods, but iteratively-reweighted methods are equipped with theoretical guarantees of convergence. In practice, as can be observed for this experiment, both methods typically display very similar behaviours in the error norm plots.

4.2. Combined regularization for spatio-temporal image deblurring.

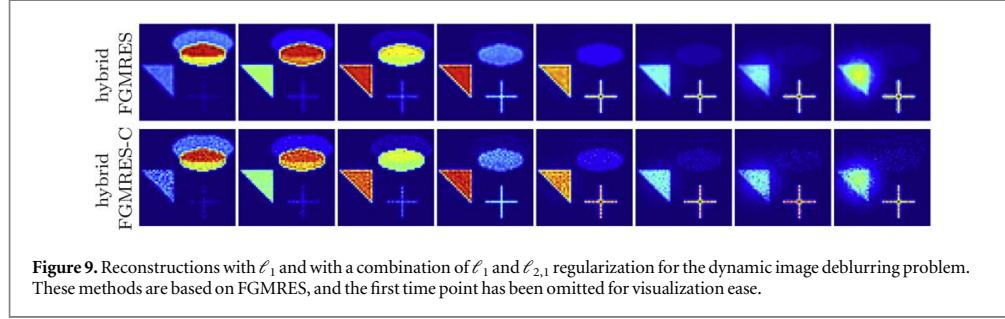
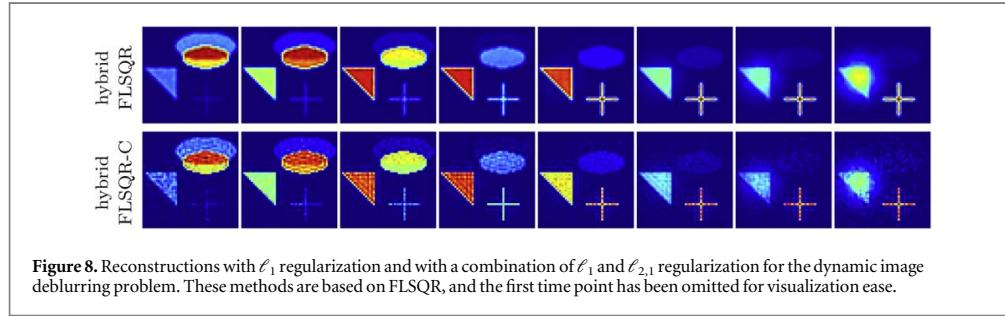
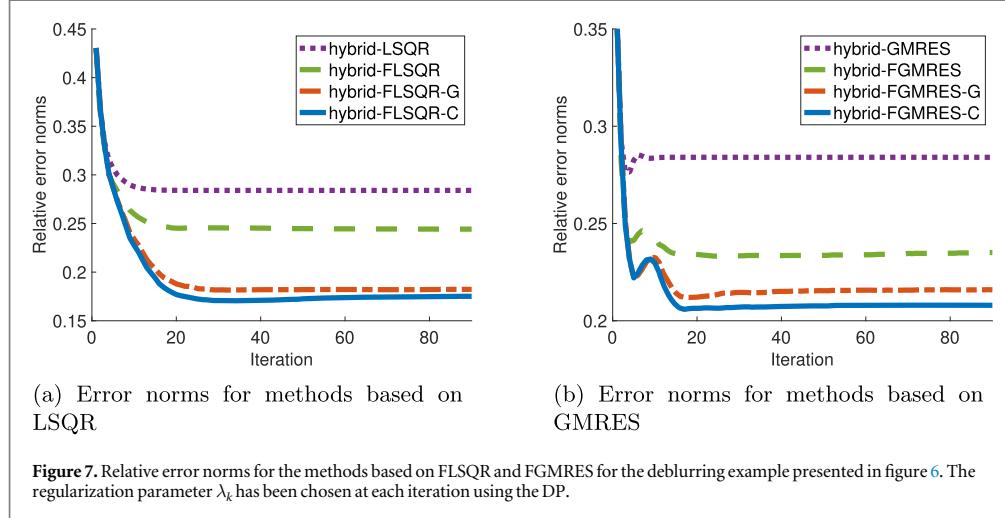
This experiment concerns a synthetic dynamic image deblurring problem, where the goal is to reconstruct a sequence of images from the sequence of their corresponding blurred and noisy counterparts [48, 49]. The blur matrix for this example is $\mathbf{A} = \mathbf{A}_t \otimes \mathbf{A}_s$ where \mathbf{A}_s represents a 2D Gaussian point spread function with spread parameter $\sigma = 1$ and bandwidth 4 and \mathbf{A}_t represents a 1D Gaussian blur with spread parameter $\sigma = 1$ and bandwidth 3. Gaussian white noise has been added to the measurements with noise level $\|\mathbf{e}\| / \|\mathbf{A}\mathbf{x}_{\text{true}}\| = 0.02$. The sequence of images forming \mathbf{x}_{true} is displayed in the first row of figure 6, while the sequence of their corresponding blurred and noisy counterparts \mathbf{b} can be found in the second row. Note that \mathbf{x} and \mathbf{b} are spatio-temporal images of size $50 \times 50 \times 9$ (i.e. 50×50 pixels at 9 time points).



In this example, some regions in all of the images in the sequence are always zero while other regions have pixels of changing intensity. It is then appropriate to use group sparsity regularization where each spatial location, over all time points, constitutes a group. We compare the results for ℓ_2 , ℓ_1 and $\ell_{2,1}$ regularization using hybrid flexible methods, and we present results for the new method combining an ℓ_1 and an $\ell_{2,1}$ regularization term. Since the forward model matrix \mathbf{A} is square, methods based on both LSQR and GMRES are tested. The error norm plots for all methods can be found in figure 7. All of the presented approaches used the DP as defined in (27) to select the regularization parameter, where \mathbf{x}_k in (27) corresponds to the approximated solution of (16) computed at iteration k . Note that τ_λ in (15) must be set ahead of the iterations. In particular, $\tau_\lambda = 1.2$ for methods based on FLSQR and $\tau_\lambda = 0.8$ for methods based on FGMRES.

It can be observed from the plot of relative reconstruction error norms, see figure 7, that ℓ_1 regularization (hybrid-FLSQR and hybrid-FGMRES) outperforms Tikhonov regularization (hybrid-LSQR and hybrid-GMRES) for this problem. Moreover, the new algorithms enforcing group sparsity regularization (hybrid-FLSQR-G and hybrid-FGMRES-G) produce significantly better reconstructions (i.e., smaller reconstruction errors), with the combined approaches (hybrid-FLSQR-C and hybrid-FGMRES-C) performing marginally better than group sparsity regularization on its own.

One can also observe that for this example the methods based on FLSQR (Figure 7(a)) perform better than the methods based on FGMRES (figure 7(b)), at the computational cost of an extra matrix-vector-product with the adjoint \mathbf{A}^\top per iteration. Nevertheless, we provide results for both solvers to highlight the benefit of using group sparsity (and a combination of sparsity and group sparsity) with respect to other regularizers, and we reiterate that this can be done using different flexible Krylov methods seamlessly. Reconstructed images corresponding to ℓ_1 regularization (top row) and a combination of ℓ_1 and $\ell_{2,1}$ regularization (bottom row) are presented for methods based on LSQR in figure 8 and for methods based on GMRES in figure 9. In particular, we would like to draw attention to the cross in the bottom-right of the images at the later time points (images on the far right). One can see that the reconstructions using ℓ_1 regularization (top row) are more blurred than the ones using a combination of ℓ_1 and $\ell_{2,1}$ (bottom row), which appear much crisper and closer to the true solution in figure 6 (top row) both for methods based on FLSQR and on FGMRES.



4.3. Anomaly detection in atmospheric inverse modeling

In this experiment, we are interested in the inverse problem of efficiently and accurately detecting anomalies from observed data in atmospheric inverse modeling. This is relevant, for example, in large-scale anomalous emissions of greenhouse gasses and air pollution detection. A major challenge for anomaly detection is that inverse models using standard priors are not able to capture both anomalies (e.g. natural gas blowouts) and smooth regions (e.g. broad-scale emissions) simultaneously. A new solution decomposition framework was described in [9] where the desired parameters are represented as a sum of different stochastic components, and different priors are used for each component. In particular, a Gaussian prior was used for the smooth component of the reconstruction and a sparsity prior was used for capturing anomalies. However, anomalies are typically sparsely distributed in space but consistent over time, and the previous framework cannot capture such phenomena. Thus, we have extended the solution decomposition framework for group sparsity regularization and present the results here.

In this example we consider a realistic linear atmospheric inverse model, where the aim is to estimate CO₂ fluxes across North America at 3-hourly temporal resolution over 41 days (approximately 6 weeks from late June

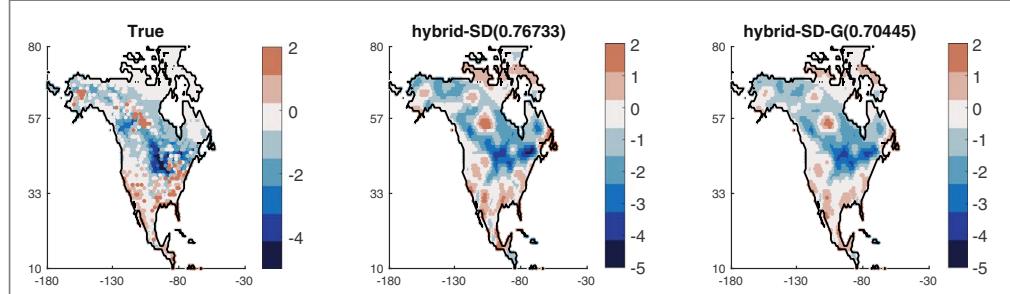


Figure 10. Anomaly detection OCO-2 example. We provide the averaged true fluxes for reference along with the averaged computed reconstruction of the fluxes (in $\mu\text{mol m}^{-2}\text{s}^{-1}$) for hybrid-SD and hybrid-SD-G. Relative reconstruction error norms for the averaged spatio-temporal fluxes are provided in the titles, and all results correspond to using the DP selected regularization parameters.

through July 2015) and at $1^\circ \times 1^\circ$ latitude-longitude spatial resolution. This setup corresponds to 3,222 unknowns per 3-hour time interval. In particular, $\mathbf{x}_{\text{true}} \in \mathbb{R}^{328 \times 3222}$ corresponds to the CO₂ fluxes from NOAA's CarbonTracker product (version 2019b). Synthetic satellite observations contained in $\mathbf{b} \in \mathbb{R}^{19,156}$ are generated as in (1) and mimic those from NASA's OCO-2 satellite, where \mathbf{A} simulates an atmospheric transport model and \mathbf{e} is added Gaussian noise to represent measurement errors. The components of the noise \mathbf{e} are considered to be uncorrelated, so the covariance matrix \mathbf{R} is $\sigma^2 \mathbf{I}$, where $\sigma = 1.1267$ has been chosen so that \mathbf{e} has a noise level of $\sigma \|\mathbf{e}\| / \|\mathbf{A}\mathbf{x}_{\text{true}}\| = 1$. We remark that although these noise levels seem high for classic inverse problems settings, they are realistic in real data inverse modeling studies using OCO-2 data [50, 51]. We refer the interested reader to [52, 53] for additional detail on the specifics of the problem setup.

Although a decomposition of $\mathbf{x}_{\text{true}} = \mathbf{s} + \boldsymbol{\xi}$ where $\mathbf{s}, \boldsymbol{\xi} \in \mathbb{R}^{328 \times 3222}$ is not available, we observe that, similar to actual atmospheric models, the true fluxes contain a combination of large, sparsely distributed values which correspond to anomalies (e.g. fires, anthropogenic emissions, or anomalies in biospheric fluxes) and smooth, broad regions of surface fluxes with small-scale variability. For defining the prior for $\boldsymbol{\xi}$, we follow similar approaches [52, 54] and consider prior covariance matrix, $\mathbf{Q} = \lambda^{-2} \mathbf{Q}_t \otimes \mathbf{Q}_s$ where \mathbf{Q}_t represents the temporal covariance and \mathbf{Q}_s represents the spatial covariance in the fluxes. These covariance matrices are defined by kernel functions

$$k_t(d_t; \theta_t) = \begin{cases} 1 - \frac{3}{2} \left(\frac{d_t}{\theta_t} \right) + \frac{1}{2} \left(\frac{d_t}{\theta_t} \right)^3 & \text{if } d_t \leq \theta_t, \\ 0 & \text{if } d_t > \theta_t, \end{cases} \quad (28)$$

$$k_s(d_s; \theta_s) = \begin{cases} 1 - \frac{3}{2} \left(\frac{d_s}{\theta_s} \right) + \frac{1}{2} \left(\frac{d_s}{\theta_s} \right)^3 & \text{if } d_s \leq \theta_s, \\ 0 & \text{if } d_s > \theta_s, \end{cases} \quad (29)$$

where d_t is day difference between two unknowns, d_s is spherical distance between two unknowns, and θ_t, θ_s are kernel parameters. In this setting, we set $\theta_t = 9.854$ and $\theta_s = 555.42$, as in [52]. For the group sparsity regularizer for \mathbf{s} , we define 3222 groups with each group corresponding to a spatial location. Thus, the groups are defined to include 41 days of 3-hourly time intervals.

We compute spatio-temporal reconstructions using hybrid-SD-G and compare the results to hybrid-SD (referred to as `sdlHybr` in [9]), where the regularization parameters λ and α are selected using the DP as defined in (27). The temporal-averaged images of the flux reconstructions are presented in figure 10, along with the temporal average of the true image. We observe that both hybrid-SD-G and hybrid-SD average reconstructions are able to capture both sources and sinks present in the true average image, with the hybrid-SD-G reconstruction having a slightly smaller relative reconstruction error for the average image. The main benefit of the solution decomposition framework is the ability to obtain two components of the solution. In figure 11, we provide the reconstructions of the individual components $\boldsymbol{\xi}$ and \mathbf{s} that form the solution for hybrid-SD (top row) and hybrid-SD-G (bottom row). It is clear that group sparsity regularization provides a smoother background and is able to distinguish persistent anomalies (as opposed to spurious false positive anomalies) better than standard sparsity regularization.

The results of this case study demonstrate that hybrid-SD-G can yield accurate results for complex, spatio-temporal atmospheric inverse modeling that are inherently different from the ones obtained using hybrid-SD. Comparisons of hybrid-SD to existing reconstruction methods can be found in [9].

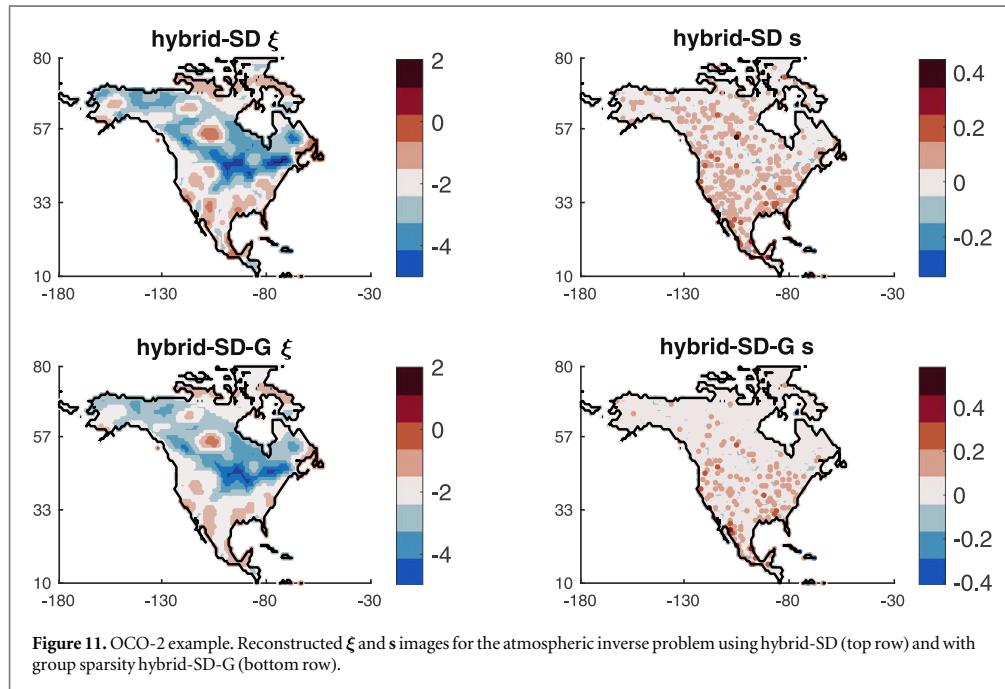


Figure 11. OCO-2 example. Reconstructed ξ and s images for the atmospheric inverse problem using hybrid-SD (top row) and with group sparsity hybrid-SD-G (bottom row).

5. Conclusions and future work

This paper presents a suite of numerical algorithms based on flexible Krylov methods for solving linear inverse problems with group sparsity regularization. The idea of using group structure to achieve better sparse recovery has received much attention, but automated and efficient approaches for practical implementation are still lacking. This work addressed that gap by proposing new iterative methods that are efficient, since they only require matrix-vector and possibly matrix-transpose-vector products, and automatic, by selecting regularization parameters on-the-fly. Moreover, these approaches exploit flexible preconditioning techniques to avoid inner-outer schemes by building a single solution subspace from which to compute solutions.

Additional regularization can be added to the projected problem to avoid semi-convergence or to guarantee convergence of the proposed method (at a higher computational cost). This scheme allows for the regularization parameter to be chosen adaptively using suitable parameter choice criteria. In particular, the results in this paper are shown for hybrid-FGMRES, hybrid-FLSQR, IRW-FLSQR, and hybrid-SD, all with group sparsity regularization and using the DP to find appropriate regularization parameters.

Moreover, the scheme presented in this paper is highly general, since very different regularization terms can be modeled in the group sparsity framework. This is highlighted in the numerical examples, which show the performance of the method in different applications and for different group sparsity modalities. In particular, an example of overlapping group sparsity is shown using the natural tree-structure of wavelet decompositions and non-overlapping group sparsity is tested in solutions that have spatio-temporal components where the solution is temporally persistent and sparse in space.

Future work includes extensions to other applications and other constraints (e.g. nonnegativity). For example, group sparsity regularization has been used for nonnegative matrix factorization [55] and for group-based dictionaries in neuroimaging using fMRI [56].

Acknowledgments

The authors gratefully acknowledge support from the Cambridge Mathematics of Information in Healthcare Hub (CMIH), University of Cambridge. The authors would also like to thank Fred Vickers-Hastings and Pol del Aguila Pla for early discussions on this work.

This work was partially supported by the National Science Foundation program under grants DMS-2341843, DMS-2245192, and DMS-2208294. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Data availability statement

The data cannot be made publicly available upon publication because no suitable repository exists for hosting data in this field of study. The data that support the findings of this study are available upon reasonable request from the authors.

Statements and declarations

Both authors contributed equally to this work, so authors are listed alphabetically.

Conflict of interest

The authors have no competing interests to declare that are relevant to the content of this article.

ORCID iDs

Julianne Chung  <https://orcid.org/0000-0002-6760-4736>

Malena Sabaté Landman  <https://orcid.org/0000-0003-3215-2773>

References

- [1] Hansen P C 2010 *Discrete Inverse Problems: Insight and Algorithms*. ed N J Higham (SIAM) (<https://doi.org/10.1137/1.9780898718836>)
- [2] Vogel C R 2002 *Computational Methods for Inverse Problems* (SIAM) (<https://doi.org/10.1137/1.9780898717570>)
- [3] Scardapane S, Comminiello D, Hussain A and Uncini A 2017 Group sparse regularization for deep neural networks *Neurocomputing* **241** 81–9
- [4] Pasha M, Saibaba A K, Gazzola S, Espanol M I and De Sturler E 2023 A computational framework for edge-preserving regularization in dynamic inverse problems *Electronic Transactions on Numerical Analysis (ETNA)* **58** pp 486–516
- [5] Huang Y and Jia Z 2017 Some results on the regularization of LSQR for large-scale discrete ill-posed problems *Science China Mathematics* **60** 701–18
- [6] Jensen T K and Hansen P C 2007 Iterative regularization with minimum-residual methods *BIT Numerical Mathematics* **47** 103–20
- [7] Chartrand R and Wohlberg B 2013 *2013 IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE nonconvex ADMM algorithm for group sparsity with sparse groups* pp 6009–13
- [8] Kaipio J P and Somersalo E 2002 Estimating anomalies from indirect observations *J. Comput. Phys.* **181** 398–406
- [9] Chung J, Jiang J, Miller S M and Saibaba A K 2024 Hybrid projection methods for solution decomposition in large-scale Bayesian Inverse Problems *SIAM Journal on Scientific Computing* **46** S97–S119
- [10] Gazzola S and Nagy J G 2014 Generalized Arnoldi-Tikhonov method for sparse reconstruction *SIAM J. Sci. Comput.* **36** 225–47
- [11] Chung J and Gazzola S 2019 Flexible Krylov methods for ℓ_p regularization *SIAM J. Sci. Comput.* **41** 149–71
- [12] Gazzola S and Sabaté Landman M 2020 Krylov methods for inverse problems: Surveying classical, and introducing new, algorithmic approaches *GAMM-Mitteilungen* **43** 202000017
- [13] Gazzola S and Sabaté Landman M 2019 Flexible GMRES for total variation regularization *BIT Numerical Mathematics* **59** 721–46
- [14] Zou H and Hastie T 2005 Regularization and variable selection via the elastic net *Journal of the royal statistical society: series B (statistical methodology)* **67** 301–20
- [15] Aguilera Pla P and Jaldén J 2018 Cell detection by functional inverse diffusion and non-negative group sparsity part ii: Proximal optimization and performance evaluation *IEEE Trans. Signal Process.* **66** 5422–37
- [16] Costa F, Batatia H, Oberlin T, D'Giano C, Tourneret J-Y and Bayesian E E G 2017 Source localization using a structured sparsity prior *NeuroImage* **144** 142–52
- [17] Majumdar A and Ward R K 2009 *2009 IEEE International Conference on Acoustics, Speech and Signal ProcessingClassification via group sparsity promoting regularization* pp 861–4
- [18] Rao N S, Nowak R D, Wright S J and Kingsbury N G 2011 *2011 XVIII IEEE International Conference on Image ProcessingConvex approaches to model wavelet sparsity patterns* pp 1917–20
- [19] Baraniuk R G, Cevher M F D V and Hedge C 2010 Model-based compressive sensing *IEEE Transactions on Information Theory* **56** 1982–2001
- [20] Xiong Z, Guleryuz O G and Orchard M T 1996 A DCT-based embedded image coder *IEEE Signal Process Lett.* **3** 289–90
- [21] Chen Chen L H and Junzhou Huang L H 2014 Fast iteratively reweighted least squares algorithms for analysis-based sparsity reconstruction *Medical Image Analysis* **49** 141–52
- [22] Chen C and Huang J 2012 Compressive sensing MRI with wavelet tree sparsity *Proceedings of the 25th International Conference on Neural Information Processing Systems* Volume 1 (Red Hook, NY, USA: Curran Associates Inc.)
- [23] Gazzola S, Nagy J G and Landman M S 2021 Iteratively reweighted FGMRES and FLSQR for sparse reconstruction *SIAM J. Sci. Comput.* **43** 47–69
- [24] Beck A and Teboulle M 2009 A fast iterative shrinkage-thresholding algorithm for linear inverse problems *SIAM J. Imaging Sci.* **2** 183–202
- [25] Wright S J, Nowak R D and Figueiredo M A T 2008 *2008 IEEE International Conference on Acoustics, Speech and Signal ProcessingSparse reconstruction by separable approximation* pp 3373–6
- [26] Daubechies I, DeVore R, Fornasier M and Güntürk C S 2010 Iteratively reweighted least squares minimization for sparse recovery *Commun. Pure Appl. Math.* **63** 1–38

[27] Rodriguez P and Wohlberg B 2008 *Proceedings of IV IEEE Andean Technical Conference (ANDESCON), Cusco, Perú* An efficient algorithm for sparse representations with ℓ_p data fidelity term

[28] Gazzola S, Scott S J and Spence A 2021 Flexible Krylov methods for edge enhancement in imaging *Journal of Imaging* **7**

[29] Beck A and Teboulle M 2009 Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems *IEEE Trans. Image Process.* **18** 2419–34

[30] Chan T F, Golub G H and Mulet P 1999 A nonlinear primal-dual method for total variation-based image restoration *SIAM J. Sci. Comput.* **20** 1964–77

[31] Osher S, Burger M, Goldfarb D, Xu J and Yin W 2005 An iterative regularization method for total variation-based image restoration *Multiscale Modeling & Simulation* **4** 460–89

[32] Rodriguez P and Wohlberg B 2009 Efficient minimization method for a generalized total variation functional *IEEE Trans. Image Process.* **18** 322–32

[33] Bach F et al 2012 Optimization with sparsity-inducing penalties *Foundations and Trends in Machine Learning* **4** 1–106

[34] Herzog R, Stadler G and Wachsmuth G 2012 Directional sparsity in optimal control of partial differential equations *SIAM J. Control Optim.* **50** 943–63

[35] Huang G, Lanza A, Morigi S, Reichel L and Sgallari F 2017 Majorization-minimization generalized Krylov subspace methods for ℓ_p - ℓ_q optimization applied to image restoration *BIT Numerical Mathematics* **57** 351–78

[36] Cho T, Chung J, Miller S M and Saibaba A K 2022 Computationally efficient methods for large-scale atmospheric inverse modeling *Geoscientific Model Development* **15** 5547–65

[37] Cho T, Chung J and Jiang J 2021 Hybrid projection methods for large-scale inverse problems with mixed Gaussian priors *Inverse Prob.* **37** 044002

[38] Calvetti D 2007 Preconditioned iterative methods for linear discrete ill-posed problems from a Bayesian inversion perspective *J. Comput. Appl. Math.* **198** 378–95

[39] Simoncini V and Szyld D B 2007 Recent computational developments in Krylov subspace methods for linear systems *Numer. Linear Algebra Appl.* **14** 1–59

[40] Chung J and Gazzola S 2024 Computational methods for large-scale inverse problems: a survey on hybrid projection methods *SIAM Review* **66** 205–84

[41] Morozov V A 1966 On the solution of functional equations by the method of regularization *Doklady Mathematics* **7** 414–7

[42] Fornasier M, Naumova V and Pereverzev S V 2014 Parameter choice strategies for multipenalty regularization *SIAM J. Numer. Anal.* **52** 1770–94

[43] Lu S, Pereverzev S V, Shao Y and Tautenhahn U 2010 Discrepancy curves for multi-parameter regularization *Journal of Inverse and Ill-Posed Problems* **18** 655–76

[44] Gazzola S and Reichel L 2016 A new framework for multi-parameter regularization *BIT Numerical Mathematics* **56** 919–49

[45] Gazzola S and Sabaté Landman M 2020 Krylov methods for inverse problems: Surveying classical, and introducing new, algorithmic approaches *GAMM-Mitteilungen* **43** 202000017

[46] Daubechies I 1992 *Ten Lectures on Wavelets* (SIAM) (<https://doi.org/10.1137/1.9781611970104>)

[47] Gazzola S, Hansen P C and Nagy J G 2018 IR Tools: a MATLAB package of iterative regularization methods and large-scale test problems *Numer. Algorithms* **81** 1–39

[48] Hansen P C 1994 Regularization tools: A MATLAB package for analysis and solution of discrete ill-posed problems *Numer. Algorithms* **6** 1–35

[49] Chung J, Saibaba A K, Brown M and Westman E 2018 Efficient generalized Golub-Kahan based methods for dynamic inverse problems *Inverse Prob.* **34** 024005

[50] Miller S M, Michalak A M, Yadav V and Tadić J M 2018 Characterizing biospheric carbon balance using CO₂ observations from the OCO-2 satellite *Atmos. Chem. Phys.* **18** 6785–99

[51] Miller S M and Michalak A M 2020 The impact of improved satellite retrievals on estimates of biospheric carbon balance *Atmos. Chem. Phys.* **20** 323–31

[52] Miller S M, Saibaba A K, Trudeau M E, Mountain M E and Andrews A E 2020 Geostatistical inverse modeling with very large datasets: an example from the orbiting carbon observatory 2 (OCO-2) satellite *Geoscientific Model Development* **13** 1771–85

[53] Liu X, Weinbren A L, Chang H, Tadić J, Mountain M E, Trudeau M E, Andrews A E, Chen Z and Miller S M 2020 Data reduction for inverse modeling: an adaptive approach v1.0 *Geoscientific Model Development Discussions* **2020** 1–21

[54] Yadav V, Michalak A M, Ray J and Shiga Y P 2016 A statistical approach for isolating fossil fuel emissions in atmospheric inverse problems *Journal of Geophysical Research: Atmospheres* **121** 12–490

[55] Kim J, Monteiro R D and Park H 2012 *Proceedings of the 2012 SIAM International Conference on Data Mining* Group sparsity in nonnegative matrix factorization pp 851–62

[56] Zhao Q, Li W X, Jiang X, Lv J, Lu J and Liu T 2018 Functional brain networks reconstruction using group sparsity-regularized learning *Brain Imaging and Behavior* **12** 758–70