



Technical Report No. 194

**Fast algorithms for
total-variation
based optimization**

Álvaro Barbero^{1,2} and Suvrit Sra¹

26 August, 2010

This Technical Report has been approved by:

Director at MPIK

Postdoc at MPIK



Technical Report No. 194

**Fast algorithms for
total-variation
based optimization**

Álvaro Barbero^{1,2} and Suvrit Sra¹

26 August, 2010

¹ MPI für biologische Kybernetik

² Universidad Autónoma de Madrid and Instituto de Ingeniería del Conocimiento

Fast algorithms for total-variation based optimization

Álvaro Barbero and Suvrit Sra

Abstract. We derive a number of methods to solve efficiently simple optimization problems subject to a total-variation (TV) regularization, under different norms of the TV operator and both for the case of 1-dimensional and 2-dimensional data. In spite of the non-smooth, non-separable nature of the TV terms considered, we show that a dual formulation with strong structure can be derived. Taking advantage of this structure we develop adaptations of existing algorithms from the optimization literature, resulting in efficient methods for the problem at hand. Experimental results show that for 1-dimensional data the proposed methods achieve convergence within good accuracy levels in practically linear time, both for L1 and L2 norms. For the more challenging 2-dimensional case a performance of order $O(N^2 \log_2 N)$ for $N \times N$ inputs is achieved when using the L2 norm. A final section suggests possible extensions and lines of further work.

1 Introduction

The total-variation (TV) operator usually appears in the context of image and signal processing problems, and can be defined as follows. Given a signal defined as a function $f : \Omega \rightarrow \mathcal{R}$, where Ω is a bounded open subset of \mathcal{R}^D , the total variation of the signal is defined as

$$TV_d(f) = \int_{\Omega} \|\nabla f(x)\|_p dx,$$

with $\|\nabla f(x)\|_p$ the L_p norm ($p \geq 1$) of the gradient of f at x [13]. If Ω is a discrete space which can be represented as an $N \times N$ matrix \mathbf{X} (an square image), this formula can be simplified down to

$$\begin{aligned} TV_2(\mathbf{X}) &= \sum_{i,j} \|\partial^+ \mathbf{X}_{i,j}\|_p, \\ \partial^+ \mathbf{X}_{i,j} &= (\mathbf{X}_{i+1,j} - \mathbf{X}_{i,j}, \mathbf{X}_{i,j+1} - \mathbf{X}_{i,j}), \end{aligned} \tag{1.1}$$

i.e. it is the norm of the discrete gradients of the image. Similarly, for a 1-dimensional signal the TV is defined as

$$TV_1(\mathbf{x}) = \sum_i \|\mathbf{x}_{i+1} - \mathbf{x}_i\|_p = \sum_i |\mathbf{x}_{i+1} - \mathbf{x}_i| \tag{1.2}$$

where the norm gets reduced to the absolute value (or L1 norm), as each discrete gradient consists of a single-entry vector.

One common application of the TV operator is the field of image denoising, where it was introduced by Rudin, Osher and Fatemi in what became to be known as the ROF model [23]. In this setting only a noisy version y of a true image x is known, and we strive to recover y from x under the assumptions that x should be similar to y and that the true image x should be smooth, i.e. the gradient of x should be small. In the ROF model this is translated to the following optimization problem

$$\min_x \frac{1}{2} \|x - y\|_2^2 + \lambda TV_d(x), \tag{1.3}$$

where λ is a penalty parameter for non-smoothness.

In this work we will show how problems in the form (1.3) can be solved efficiently via a transformation to a simpler dual formulation, which allows to take advantage of the structure induced by the total variation norm. We will then analyze several specific cases of this formulation, and show how standard optimization algorithms can be

where $\text{vec}(\mathbf{X})$ returns a vectorized version of \mathbf{X} , constructed by stacking the columns of \mathbf{X} . This results in the optimization problem

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{B} \text{vec}(\mathbf{X})\|_p. \quad (2.5)$$

Again, this formulation is equivalent to (1.1) only when $p = 1$. Conversely, these expressions give rise to different models for any $p > 1$. These differences can be stated more clearly when interpreting them as mixed norms. Recall that a mixed norm is defined as

$$\|\mathbf{x}\|_{p,q} = \left(\sum_a \left(\sum_b |x_{ab}|^p \right)^{q/p} \right)^{1/q},$$

which essentially boils down to considering a block structure in x , applying a norm p in each block, and then applying a norm q along blocks. Under this point of view, the original TV operator (1.1) can be regarded as a $(p, 1)$ mixed norm, in which each block is a two-entry vector component. On the other hand, our proposed approach (2.4) is nothing but a standard p norm. Though this results in a different optimization problem (and thus, model), we will see that our approach has computational advantages over the more difficult to handle mixed norm.

Whatever the case, we see that the TV operator can be rewritten (or approximated) as an unconstrained convex problem in the form

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{M}\mathbf{x}\|_p. \quad (2.6)$$

While the first term is differentiable, the second term is non-smooth and, more importantly, non-separable, i.e. cannot be broken down into independent components involving single entries x_i . This kind of problem has been addressed thoroughly in the literature in the form of the so-called ‘‘proximity operators’’ [8], which have the form

$$\text{Prox}_R(y) = \arg \min_{x \in X} \frac{1}{2} \|x - y\|_2^2 + R(x), \quad (2.7)$$

where R is another operator. Thus, problem (2.6) is nothing but the proximity operator of the total variation. We will later see in more detail the importance of this connection. For now, the interesting point is the fact that if we had a matrix $\mathbf{M} = \mathbf{I}$ then (2.6) would turn out to be separable. What is more, proximity operators for these cases exist, which allow to find a solution in closed form for the $p = 1$ and $p = 2$ norm cases, and in roughly linear time for the $p = \infty$ norm [16]. This observation will also be of use in our proposed methods.

2.2 Duality of the TV problem

We show now how the proximity problem (2.6) can be casted into a simple dual problem, and how we can recover the solution of (2.6) from the solution of this new problem. This can be done by applying Theorem 31.2 in [21], though we present here an alternative, more straightforward proof. To do so, we first perform the derivations for a more general form as follows.

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a twice-differentiable, strictly convex function; let $r : \mathbb{R}^m \rightarrow \mathbb{R}$ be a closed, proper convex function, and let $\mathbf{B} \in \mathbb{R}^{m \times n}$ be a matrix (can be a linear operator in general, but we ignore that for now). Consider the optimization problem

$$\min_{\mathbf{x}} f(\mathbf{x}) + \lambda r(\mathbf{B}\mathbf{x}), \quad (2.8)$$

where $\lambda > 0$ is a scalar. To derive the dual of (2.8) introduce an auxiliary variable $\mathbf{z} = \mathbf{B}\mathbf{x}$, whereby (2.8) becomes

$$\min_{\mathbf{x}, \mathbf{z}} f(\mathbf{x}) + \lambda r(\mathbf{z}), \quad \text{s.t. } \mathbf{z} = \mathbf{B}\mathbf{x}. \quad (2.9)$$

Let \mathbf{u} be the dual-variables corresponding to the equality constraint. Then, the Lagrangian to (2.9) is

$$L(\mathbf{x}, \mathbf{z}, \mathbf{w}) = f(\mathbf{x}) + \lambda r(\mathbf{z}) + \mathbf{w}^T (\mathbf{B}\mathbf{x} - \mathbf{z}). \quad (2.10)$$

The dual function $g(\mathbf{w}) = \inf_{\mathbf{x}, \mathbf{z}} L(\mathbf{x}, \mathbf{z}, \mathbf{w})$, so that

$$g(\mathbf{w}) = \left(\inf_{\mathbf{x}} f(\mathbf{x}) + \mathbf{w}^T \mathbf{B}\mathbf{x} \right) + \left(\inf_{\mathbf{z}} -\mathbf{w}^T \mathbf{z} + \lambda r(\mathbf{z}) \right). \quad (2.11)$$

Both infima in (2.11) may be written as suprema, so that

$$g(\mathbf{w}) = \left(-\sup_{\mathbf{x}} -\mathbf{w}^T \mathbf{B}\mathbf{x} - f(\mathbf{x}) \right) - \left(\sup_{\mathbf{z}} \mathbf{w}^T \mathbf{z} - \lambda r(\mathbf{z}) \right), \quad (2.12)$$

which are nothing but the Fenchel-duals¹ of f and (λr) . Thus we obtain

$$g(\mathbf{w}) = -f^*(-\mathbf{B}^T \mathbf{w}) - \lambda r^*(\lambda^{-1} \mathbf{w}). \quad (2.13)$$

Now, in our case we have $f(\mathbf{x}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2$, and $r(\mathbf{x}) = \|\mathbf{x}\|_p$. Then $r^*(\mathbf{w}) = \delta_{\|\mathbf{w}\|_q \leq 1}$, the indicator function for dual norm q , which is defined as $\frac{1}{p} + \frac{1}{q} = 1$. Thus, (2.13) simplifies in this case to

$$g(\mathbf{w}) = \begin{cases} -\frac{1}{2} \mathbf{w}^T \mathbf{B} \mathbf{B}^T \mathbf{w} + \mathbf{w}^T \mathbf{B} \mathbf{y} & \text{where } \|\mathbf{w}\|_q \leq \lambda, \\ -\infty & \text{otherwise.} \end{cases} \quad (2.14)$$

Note: completing the square allows us to write the dual-optimization problem as

$$\min_{\mathbf{w}} \quad \frac{1}{2} \|\mathbf{B}^T \mathbf{w} - \mathbf{y}\|_2^2 \quad \text{s.t. } \|\mathbf{w}\|_q \leq \lambda. \quad (2.15)$$

Which is, in fact, an instance of the so-called Trust Region Subproblem, in which we optimize a local quadratic model of some function subject to not leaving a trust region defined by an L_q -norm. This kind of problems have been studied in depth in the literature, and as we shall see, standard algorithms exist to find a solution efficiently.

Once the solution \mathbf{w}^* to (2.15) has been found we can recover its related primal solution \mathbf{x}^* by noting that, by the Karush–Kuhn–Tucker optimality conditions, the gradient of the Lagrangian w.r.t. to the primal variables must be 0 at the optimum. In particular

$$\begin{aligned} \frac{\partial L(\mathbf{x}, \mathbf{z}, \mathbf{w})}{\partial \mathbf{x}} &= 0, \\ \frac{\partial}{\partial \mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \mathbf{B}^T \mathbf{w} &= 0, \\ \mathbf{x} &= \mathbf{y} - \mathbf{B}^T \mathbf{w}, \end{aligned}$$

therefore being able to recover the primal solution easily. What is more, we can also compute the dual gap with this by noting that for \mathbf{x} and \mathbf{w} feasible solutions of the primal and dual problems we have

$$\begin{aligned} \text{gap}(\mathbf{x}, \mathbf{w}) &= \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{B}\mathbf{x}\|_p - \left(-\frac{1}{2} \|\mathbf{B}^T \mathbf{w}\|_2^2 + \mathbf{w}^T \mathbf{B}\mathbf{y} \right), \\ &= \lambda \|\mathbf{B}\mathbf{x}\|_p - \mathbf{w}^T \mathbf{B}\mathbf{y} + \|\mathbf{B}^T \mathbf{w}\|_2^2, \\ &= \lambda \|\mathbf{B}\mathbf{x}\|_p - \mathbf{w}^T \mathbf{B}\mathbf{y} + \mathbf{w}^T \mathbf{B} \mathbf{B}^T \mathbf{w}, \\ &= \lambda \|\mathbf{B}\mathbf{x}\|_p - \mathbf{w}^T \mathbf{B}(\mathbf{y} - \mathbf{B}^T \mathbf{w}), \\ &= \lambda \|\mathbf{B}\mathbf{x}\|_p + \mathbf{w}^T (-\mathbf{B}\mathbf{x}), \end{aligned}$$

which is fast to compute once $-\mathbf{B}\mathbf{x}$ is known. The dual gap is a good measure of the quality of an approximate solution $\tilde{\mathbf{w}}$ of the dual, thus being useful as a stopping criterion for an optimization algorithm.

2.3 Structure induced by the TV operator

The special structure of matrices \mathbf{D} and \mathbf{B} allows to perform very fast products and factorizations, not only because of their sparsity, but also because of the way their non-zero entries are distributed. Taking advantage of these facts will prove to be vital for the implementation of an efficient algorithm.

In the 1-dimensional case the matrix \mathbf{D} presents itself as a bidiagonal matrix, which allows to compute matrix-vector products in linear time as follows

$$\mathbf{D}\mathbf{v} = \mathbf{v}_{2:N} - \mathbf{v}_{1:N-1}, \quad \mathbf{D}^T \mathbf{v} = \begin{bmatrix} 0 \\ \mathbf{v} \end{bmatrix} - \begin{bmatrix} \mathbf{v} \\ 0 \end{bmatrix},$$

¹Fenchel-dual is defined as: $f^*(\mathbf{w}) = \sup_{\mathbf{x}} \mathbf{w}^T \mathbf{x} - f(\mathbf{x})$

where $\mathbf{v}_{i:j}$ denotes the vector formed by the entries \mathbf{v}_i to \mathbf{v}_j (like in MATLAB notation). Furthermore, the matrices $\mathbf{D}\mathbf{D}^T$ and $\mathbf{D}^T\mathbf{D}$ present a tridiagonal structure, which allow for an efficient Cholesky-like factorization in the form $\mathbf{M} = \mathbf{L}\mathbf{\Delta}\mathbf{L}^T$. Also methods for fast eigendecomposition for $\mathbf{D}\mathbf{D}^T$ are available. We will make use of these fast factorizations in our proposed algorithms. Technical details about how to perform them are presented in Appendix A.

Regarding the matrix \mathbf{B} used in the 2-dimensional case, it can be nicely rewritten in terms of \mathbf{D} and the Kronecker product operator as

$$\mathbf{B} = \begin{pmatrix} \mathbf{I}_N & \otimes & \mathbf{D} \\ \mathbf{D} & \otimes & \mathbf{I}_N \end{pmatrix}, \quad (2.16)$$

with \mathbf{I}_N the $N \times N$ identity matrix. The Kronecker product hence allows to express \mathbf{B} compactly, while featuring also a series of nice properties [24]. In our application the most useful of these properties is

$$\mathbf{Y} = \mathbf{C}\mathbf{X}\mathbf{B}^T \equiv \mathbf{y} = (\mathbf{B} \otimes \mathbf{C})\mathbf{x}, \quad (2.17)$$

with $\mathbf{y} = \text{vec}(\mathbf{Y})$ and $\mathbf{x} = \text{vec}(\mathbf{X})$. In other words, we are able to transform an equation system expressed through Kronecker products into a matrix equation system. This is of immediate use for computing matrix-vector products involving \mathbf{B} , as

$$\mathbf{B}\mathbf{v} = \begin{bmatrix} \mathbf{D}\mathbf{V} \\ \mathbf{V}\mathbf{D}^T \end{bmatrix}, \quad \mathbf{B}^T\mathbf{w} = \mathbf{D}^T\mathbf{W}_1 + \mathbf{W}_2\mathbf{D}, \quad (2.18)$$

with $\mathbf{v} \in \mathbb{R}^{N^2}$, $\mathbf{w} \in \mathbb{R}^{2(N-1)N}$, $\mathbf{v} = \text{vec}(\mathbf{V}) \in \mathbb{R}^{N \times N}$, $\mathbf{w} = \begin{bmatrix} \text{vec}(\mathbf{W}_1) \\ \text{vec}(\mathbf{W}_2) \end{bmatrix}$, $\mathbf{W}_1 \in \mathbb{R}^{(N-1) \times N}$, $\mathbf{W}_2 \in \mathbb{R}^{N \times (N-1)}$. This saves the need for maintaining an explicit representation for \mathbf{B} in memory.

We shall use these structures properties to guarantee fast computations in the methods we will introduce next.

3 1-dimensional Total Variation solvers

3.1 Alternating-Direction Method of Multipliers solver

The primal problem can be solved directly by using the so-called Alternating-Direction Method of Multipliers (ADMM), [8]. This method works on objective functions in the form

$$\min_{\mathbf{x}, \mathbf{b}} f(\mathbf{x}) + g(\mathbf{b}) \quad \text{s.t.} \quad \mathbf{L}\mathbf{x} = \mathbf{b},$$

which is our case with $f(\mathbf{x}) = \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|_2^2$, $g(\mathbf{b}) = \lambda\|\mathbf{b}\|_p$ and $\mathbf{L} = \mathbf{D}$. The method proceeds by alternatively computing the proximity operator of f and g , eventually converging to an optimum of the problem. The general procedure is shown in Algorithm 1, where $\text{prox}_f^L(\mathbf{y}) = \arg \min_{\mathbf{x}} f(\mathbf{x}) + \frac{1}{2}\|\mathbf{L}\mathbf{x} - \mathbf{y}\|_2^2$.

Algorithm 1 Alternating-Direction Method of Multipliers

Inputs: functions $f(\mathbf{x})$, $g(\mathbf{b})$, matrix \mathbf{L} , augmented Lagrangian index $\gamma > 0$, stopping tolerance τ .

Initialization: set initial values for \mathbf{p} , \mathbf{z} .

while stopping criterion $> \tau$ **do**

$\mathbf{x} = \text{prox}_{\gamma f}^L(\mathbf{p} - \mathbf{z})$

$\mathbf{s} = \mathbf{L}\mathbf{x}$

$\mathbf{p} = \text{prox}_{\gamma g}(\mathbf{s} + \mathbf{z})$

$\mathbf{z} = \mathbf{z} + \mathbf{s} - \mathbf{p}$

end while

The speed of the algorithm depends heavily on how efficiently can we solve the proximity operators for f and g . Luckily, in the problem at hand they result to be solvable in closed form or at a linear cost, as on the one hand

$$\begin{aligned}
\text{prox}_{\gamma f}^L(\mathbf{p} - \mathbf{z}) &= \arg \min_{\mathbf{x}} \gamma f(\mathbf{x}) + \frac{1}{2} \|\mathbf{L}\mathbf{x} - \mathbf{p} + \mathbf{z}\|_2^2, \\
&= \arg \min_{\mathbf{x}} \gamma \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \frac{1}{2} \|\mathbf{L}\mathbf{x} - \mathbf{p} + \mathbf{z}\|_2^2. \\
0 &= \frac{\partial}{\partial \mathbf{x}}, \\
&= \gamma(\mathbf{x} - \mathbf{y}) + \mathbf{D}^T(\mathbf{D}\mathbf{x} - \mathbf{p} + \mathbf{z}), \\
(\gamma \mathbf{I} + \mathbf{D}^T \mathbf{D})\mathbf{x} &= \gamma \mathbf{y} + \mathbf{D}^T(\mathbf{p} - \mathbf{z}),
\end{aligned}$$

and so the proximity operator can be computed by solving the tridiagonal equation system defined by $\gamma \mathbf{I} + \mathbf{D}^T \mathbf{D}$, which is singular for a large enough $\gamma > 0$. On the other hand we have

$$\begin{aligned}
\text{prox}_{\gamma g}(\mathbf{s} + \mathbf{z}) &= \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{s} - \mathbf{z}\|_2^2 + \gamma g(\mathbf{x}), \\
&= \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{s} - \mathbf{z}\|_2^2 + \gamma \lambda \|\mathbf{x}\|_p, \\
&= \text{prox}_{\gamma \lambda \|\cdot\|_p}(\mathbf{s} + \mathbf{z}),
\end{aligned}$$

which essentially is the proximity operator for a standard L_p norm. It has been shown that for the L1 and L2 cases this operator can be computed in closed form, while for the L_∞ case a solution can be found in linear time [8]. For other choices of p no efficient algorithm exists up to our knowledge. Therefore the ADMM method will only be useful when using the aforementioned norms.

Taking all these considerations into account, a modified algorithm for TV- L_p norm problem is shown as Algorithm 2. As a stopping criterion we measure the absolute amount of change in \mathbf{x} , as $\|\mathbf{x}^t - \mathbf{x}^{t-1}\|_1$.

Algorithm 2 Alternating-Direction Method of Multipliers for TV- L_p norm

Inputs: reference data \mathbf{y} , augmented Lagrangian index $\gamma > 0$, stopping tolerance τ , norm order p , penalty term λ .

Initialization: set initial values for \mathbf{p} , \mathbf{z} .

Precompute tridiagonal factorization of $(\gamma \mathbf{I} + \mathbf{D}^T \mathbf{D})$. (LAPACK: DPTTRF)

while stopping criterion $> \tau$ **do**

 Solve tridiagonal system $(\gamma \mathbf{I} + \mathbf{D}^T \mathbf{D})\mathbf{x} = \gamma \mathbf{y} + \mathbf{D}^T(\mathbf{p} - \mathbf{z})$. (LAPACK: DPTTRS)

$\mathbf{s} = \mathbf{D}\mathbf{x}$

$\mathbf{p} = \text{prox}_{\gamma \lambda \|\cdot\|_p}(\mathbf{s} + \mathbf{z})$

$\mathbf{z} = \mathbf{z} + \mathbf{s} - \mathbf{p}$

end while

3.2 TV-L1 case

We now introduce two algorithms to solve (2.15) efficiently for the case where $p = 1$ (TV-L1 norm). First note that the dual norm for this case is $q = \infty$, which is equivalent to impose a lower and upper bound on every entry of w ,

$$\begin{aligned}
\min_{\mathbf{w}} \quad & \frac{1}{2} \|\mathbf{D}^T \mathbf{w} - \mathbf{y}\|_2^2 \\
\text{s.t.} \quad & -\lambda \leq \mathbf{w} \leq \lambda,
\end{aligned} \tag{3.1}$$

which is a very simple box-constrained quadratic problem. Several methods are available in the literature to solve this class of problems efficiently, like TRON [18], L-BFGS-B [6], Projected Newton [4] or SBB [17]. Here we will adapt the last two ones.

$$\begin{aligned}\alpha' &= \alpha_C, \\ \beta' &= (\beta \odot e_{I-1})_I,\end{aligned}$$

where x_C is the vector formed by only the entries of x indexed by elements of C , I is the complementary set of C , $I - 1$ stands for the set of indexes of C being subtracted 1, e_C is an indicator vector with zeros everywhere except in the entries indexed by C , which are 1, and \odot is the point product.

Proof. The proof for α' is immediate, as α_i only appears in $A_{C,C}$ if $i \in C$. For β' , consider first that only one row/column i is to be removed. If $i = 1$ only β_1 does not appear in $A_{C,C}$ from the sub/superdiagonals. Same applies for $i = N$, where only β_{N-1} vanishes. If any other i applies, then both β_{i-1} and β_i are vanish, but note that in the resulting matrix the entries immediately to the right and under $\alpha_i - 1$ will be valued 0. Therefore, what is obtained is essentially the result of setting β_{i-1} to 0 and removing β_i from the resulting β' . Applying this procedure to every $i \in C$ produces the presented formula. \square

The tridiagonality property of the reduced Hessian $A_{C,C}$, allows to solve the system $A_{N,N}d_N = \nabla f(x)_N$ for d_C efficiently through the use of the tridiagonal factorization of $A_{N,N}$ (see Appendix A).

Regarding the Armijo rule, we can make use again of the fact that $x'_I = x_I$, and so the rule simplifies down to $f(x) - f([x - \beta^m d]_P) \geq \sigma \beta^m \nabla f(x)_C \cdot d_C$. We can further simplify this by defining $\delta = [x - \beta^m d]_P - x$ and noting that

$$\begin{aligned}f(x) - f(x + \delta) &= f(x) - \left(f(x) + \frac{1}{2} \delta^T A \delta^T + x^T A \delta^T + b^T \delta \right) \\ &= -\left(\frac{1}{2} \delta^T A + x^T A + b^T \right) \delta \\ &= -\left(\frac{1}{2} \delta_C^T A_{I,I} + x^T \begin{pmatrix} A_{C,C} \\ A_{I,C} \end{pmatrix} + b_C^T \right) \delta_C\end{aligned}$$

The non-zero entries of $A_{C,C}$ have already been computed, and multiplying this matrix with any vector can be done in linear time, owing again to this matrix being tridiagonal. However, $A_{I,C}$ has neither been computed, nor it is tridiagonal. Nevertheless it should be realized that its rows contain at most two nonzero entries that can be easily located. This can be seen by noting that for any $i \in I$ the row $A_{i,C}$ cannot contain the α_i term as $i \notin C$, and will only keep β_{i-1} if $i - 1 \in C$, β_i if $i + 1 \in C$. Thus, for each row of $A_{I,C}$ we can decide fastly which entries are non-zero, and are able to multiply it with a vector spending at most two product operations. It should also be realized that the term $x^T \begin{pmatrix} A_{C,C} \\ A_{I,C} \end{pmatrix}$ is independent of δ_C , and as such we only need to compute it once per iteration, regardless of the number of Armijo rule evaluations needed.

Finally, another point of improvement of the standard Projected Newton is the gradient recomputation. As in this case we have a quadratic function, the gradient can be updated as

$$\begin{aligned}\nabla f(x + \delta) &= \nabla f(x) + A \delta \\ &= \nabla f(x) + \begin{pmatrix} A_{C,C} \\ A_{I,C} \end{pmatrix} \delta_C\end{aligned}$$

where we can use similar strategies to the ones presented for the evaluation of the Armijo rule.

Summing up, the adapted Projection Newton algorithm for quadratic functions with tridiagonal Hessian is shown in Algorithm 4. As a stopping criterion we use the quantity of total violation of the Karush–Kuhn–Tucker optimality conditions, which can be shown to be

$$viol(x) = \|\nabla f(x)_C\|_1.$$

Algorithm 4 Projected Newton for tridiagonal Hessian problems

Inputs: objective function in the form $f(x) = x^T Ax + b^T x + c$, A tridiagonal defined by (α, β) , box constraints $l \leq x \leq u$, initial guess x^0 , comparison tolerance ϵ , stopping tolerance τ , Armijo parameters β, σ .

Definitions: $([x]_P)_i = \min\{\max\{x_i, l_i\}, u_i\}$, orthogonal projection of x onto feasible set.

Initialization $x = x^0$, $\nabla f(x) = Ax + b$.

while stopping criterion $> \tau$ **do**

 Identify set of active constraints

$$I = \{i | (u_i - \epsilon \leq x, \nabla f(x)_i \leq -\epsilon) \text{ or } (l_i + \epsilon \geq x, \nabla f(x)_i \geq \epsilon)\}$$

 Construct nonzero entries of reduced Hessian: $A_{C,C}$

 Compute tridiagonal factorization: $A_{C,C} = L\Delta L^T$. (LAPACK: DPTTRF)

 Compute updating direction d solving system $A_{C,C}d_C = \nabla f(x)_C$. (LAPACK: DPTTRS)

 Armijo rule: find smallest integer $m \geq 0$ such that for $\delta = [x - \beta^m d]_P - x$,

$$-\left(\frac{1}{2}\delta_C^T A_{C,C} + x^T \begin{pmatrix} A_{C,C} \\ A_{I,C} \end{pmatrix} + b_C^T\right)\delta_C \geq \sigma\beta^m \nabla f(x)_C \cdot d_C.$$

 Update position: $x' = x + \delta$.

 Update gradient: $\nabla f(x') = \nabla f(x) + \begin{pmatrix} A_{C,C} \\ A_{I,C} \end{pmatrix} \delta_C$.

end while

3.2.2 Subspace BB

The Subspace BB method [17] can be regarded as an extension of the Barzilai and Borwein method for the nonnegative constrained case [3], which combines Barzilai-Borwein stepsizes with a projected gradient approach. Being a first-order method (no Hessian information is required), the matrix D is only used through matrix-vector products, and so each iteration can be run in linear time owing to the bidiagonality of D .

An outline of the SBB algorithm adapted to the problem at hand is shown as Algorithm 5. At each iteration the set of active constraints is identified, a gradient ∇f confined to the subspace defined by the non-active constraints is constructed, and finally a BB-like stepsize is performed in the direction of the (full) gradient, projecting the result back into the feasible set.

Even though the original SBB method was designed to deal only with nonnegativity constraints, it can be easily extended to a box constrained setting. Furthermore in the original SBB algorithm every M iterations a sufficient descent check is performed, decreasing a stepsize multiplier if this check is not met; we have removed this check in our implementation, as in all the experiments performed the algorithm managed to achieve convergence without resorting to it.

Algorithm 5 Subspace BB for TV-L1

Inputs: reference data y , penalty parameter λ , initial guess x^0 , stopping tolerance τ .

Definitions: $([x]_P)_i = \min\{\max\{x_i, -\lambda\}, \lambda\}$, orthogonal projection of x onto feasible set.

Initialization $x = x^0$.

while stopping criterion $> \tau$ **do**

 Compute gradient $\nabla f = D(D^T x - y)$

 Identify set of active constraints

$$I = \{i | (-\lambda == x_i, \nabla f(x)_i \leq 0) \text{ or } (\lambda == x_i, \nabla f(x)_i \geq 0)\}$$

 Construct subspace gradient: $[\nabla \tilde{f}]_i = [\nabla f]_i$ if $i \notin I$, 0 otherwise.

if iteration number is odd **then**

$$\text{ Compute stepsize: } \alpha = \frac{\|\nabla \tilde{f}\|_2^2}{\langle \nabla \tilde{f}, DD^T \nabla \tilde{f} \rangle}.$$

else

$$\text{ Compute stepsize: } \alpha = \frac{\langle \nabla \tilde{f}, DD^T \nabla \tilde{f} \rangle}{\|DD^T \nabla \tilde{f}\|_2^2}.$$

end if

$$x = [x - \alpha \nabla f]_+.$$

end while

3.3 TV-L2 case

When $p = 2$ in (2.3) the dual problem also fancys a $q = 2$ norm, and so the feasible space is an euclidean ball of radius λ .

$$\min_{\mathbf{w}} \quad \frac{1}{2} \|\mathbf{D}^T \mathbf{w} - \mathbf{y}\|_2^2 \quad \text{s.t.} \quad \|\mathbf{w}\|_2 \leq \lambda. \quad (3.2)$$

This is the most common situation in Trust Region Subproblems, and as such a large number of solvers are available in the literature, like LSTRS [22], IP-SSM [11] and the method proposed by Moré and Sorensen [19]. A priori LSTRS is the best option, as it is a well-stablished method and good public implementations are available. However, the same LSTRS authors recommend to use Moré and Sorensen's method for the particular case in which the Cholesky decomposition of $\mathbf{A} + \delta \mathbf{I}$ can be computed efficiently, \mathbf{A} the Hessian of the objective function and $\delta \geq 0$. As we will see, this is the situation in our particular problem, and so we will adapt this method for our case.

The motivation behind Moré and Sorensen's method can be easily explained by analyzing the KKT conditions of the problem at hand. We will do this first for a general quadratic objective function in the form $\mathbf{g}^T \mathbf{w} + \frac{1}{2} \mathbf{w}^T \mathbf{A} \mathbf{w}$, as in [19]. The Lagrangian of the problem can be written as

$$L(\mathbf{w}, \alpha) = \mathbf{g}^T \mathbf{w} + \frac{1}{2} \mathbf{w}^T \mathbf{A} \mathbf{w} + \alpha (\|\mathbf{w}\|_2^2 - \lambda^2),$$

which at the optimum must meet

$$\begin{aligned} \frac{\partial L(\mathbf{w}, \alpha)}{\partial \mathbf{w}} &= 0, \\ \mathbf{g} + \mathbf{A} \mathbf{w} + \alpha \mathbf{w} &= 0, \\ (\mathbf{A} + \alpha \mathbf{I}) \mathbf{w} &= -\mathbf{g}. \end{aligned}$$

Also, as a Lagrange multiplier, $\alpha \geq 0$ and due to complementarity constraints $\alpha (\|\mathbf{w}\|_2^2 - \lambda^2) = 0$. Joining all these constraints, we can rewrite the problem as a Constraint Satisfaction Problem in the form

$$\text{find}_{\mathbf{w}} \quad \text{s.t.} \quad \begin{cases} (\mathbf{A} + \alpha \mathbf{I}) \mathbf{w} = -\mathbf{g} \\ \alpha \geq 0 \\ \alpha (\|\mathbf{w}\|_2^2 - \lambda^2) = 0 \end{cases}. \quad (3.3)$$

Two situations might happen in the solution of (3.3): either \mathbf{w} is the interior of the feasible set ($\|\mathbf{w}\|_2 \leq \lambda$) or it is in the boundary ($\|\mathbf{w}\|_2 = \lambda$). If \mathbf{w} is in the interior we necessarily have $\alpha = 0$, and then the solution is given by finding a solution to the system $\mathbf{A} \mathbf{w} = -\mathbf{g}$. The argument is also valid backwise: if the solution \mathbf{w} of the system $\mathbf{A} \mathbf{w} = -\mathbf{g}$ is in the interior of the feasible set, then \mathbf{w} solves (3.3). Thus, we can easily check whether the solution is in the interior, and if it is not, we can safely assume $\|\mathbf{w}\|_2 = \lambda$ and work only in the boundary case, which we address in the following.

The boundary condition is the key to solving the problem. To see this, suppose for now that we have α s.t. $(\mathbf{A} + \alpha \mathbf{I})$ is definite postive. Then we have

$$\begin{aligned} \lambda^2 &= \|\mathbf{w}\|_2^2 \\ &= \|-(\mathbf{A} + \alpha \mathbf{I})^{-1} \mathbf{g}\|_2^2 \\ &= \mathbf{g}^T (\mathbf{A} + \alpha \mathbf{I})^{-1T} (\mathbf{A} + \alpha \mathbf{I})^{-1} \mathbf{g} \end{aligned}$$

We can now use the eigendecomposition $\mathbf{A} = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^T$, with $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_N)$ diagonal matrix of eigenvalues sorted in increasing order and \mathbf{Q} orthogonal matrix of eigenvectors,

$$\begin{aligned}
\lambda^2 &= \mathbf{g}^T (\mathbf{Q}\Lambda\mathbf{Q}^T + \alpha\mathbf{I})^{-1T} (\mathbf{Q}\Lambda\mathbf{Q}^T + \alpha\mathbf{I})^{-1} \mathbf{g}, \\
&= \mathbf{g}^T (\mathbf{Q}(\Lambda + \alpha\mathbf{I})\mathbf{Q}^T)^{-1T} (\mathbf{Q}(\Lambda + \alpha\mathbf{I})\mathbf{Q}^T)^{-1} \mathbf{g}, \\
&= \mathbf{g}^T \mathbf{Q}(\Lambda + \alpha\mathbf{I})^{-1} \mathbf{Q}^T \mathbf{Q}(\Lambda + \alpha\mathbf{I})^{-1} \mathbf{Q}^T \mathbf{g}, \\
&= \mathbf{g}^T \mathbf{Q}(\Lambda + \alpha\mathbf{I})^{-1} (\Lambda + \alpha\mathbf{I})^{-1} \mathbf{Q}^T \mathbf{g}, \\
&= \sum_i \frac{\gamma_i^2}{(\lambda_i + \alpha)^2}, \\
&= \psi(\alpha)
\end{aligned}$$

with $\gamma_i = \mathbf{g}^T \mathbf{Q}_i$, \mathbf{Q}_i i -th column of \mathbf{Q} . Thus, the norm of the solution vector is essentially a nonlinear function ψ of α . Therefore, to solve the problem we need to find an α s.t. $\psi(\alpha) = \lambda^2$. However, the nonlinearity of ψ together with the existence of second-order poles at the points where $\alpha = -\lambda_i$ might make this task difficult. In [19] the following strategy is recommended: locate the rightmost pole, which is the one associated with the smallest eigenvalue λ_1 , and try to find an α that satisfies the constraint in the ‘‘poles-free’’ interval $\alpha \in (-\lambda_1, \infty)$. This approach is reasonable, as $\lim_{\alpha \rightarrow -\lambda_1} = \infty$, $\lim_{\alpha \rightarrow \infty} = 0$ and by continuity $\exists \alpha^*$ s.t. $\psi(\alpha^*) = \lambda^2$. Unfortunately, this argument might fail, in particular if $\gamma_1 = 0$ then α^* might be located in the interval $(-\infty, -\lambda_1]$. This is the so-called ‘‘hard-case’’, and to deal with it additional safeguarding measures need to be taken in Moré and Sorensen’s algorithm, also resulting in slower convergence. Fortunately, in the problem at hand the hard case cannot take place, namely,

Proposition 2. *In a problem in the form (3.2), Moré and Sorensen’s ‘‘hard case’’ cannot happen.*

Proof. The rows of the Hessian of (3.2), $\mathbf{A} = \mathbf{D}\mathbf{D}^T$, are linearly independent, thus rendering \mathbf{A} positive definite. As such, every one of its eigenvalues $\lambda_i > 0$. Consequently, the rightmost pole is located at $-\lambda_1 < 0$, and as by the KKT conditions (3.3) $\alpha \geq 0$, then the solution $\alpha^* \in [0, \infty)$. \square

Corollary 3. *It follows that the maximum value of the $\psi(\alpha)$ function is attained at $\alpha = 0$. Therefore, for any $\lambda > \lambda_{MAX} = \psi(0)$ the solution \mathbf{w} is always the same point in the interior of the feasible set. Thus, the range of λ values generating different models is bounded by $\lambda \in [0, \lambda_{MAX}]$.*

Therefore, we do not need to worry about the hard case, and can just look for α^* in the pole-free, smooth interval $[0, \infty)$. To do so, Moré and Sorensen recommend to use Newton’s method for locating the roots of the function

$$\phi(\alpha) = \frac{1}{\lambda} - \frac{1}{\|\mathbf{w}\|_2} = 0,$$

which is almost linear in the search interval, guaranteeing very fast convergence. Recall that Newton’s method involves updating our estimate of the function’s root as $\alpha' = \alpha - \frac{\phi(\alpha)}{\phi'(\alpha)}$ with $\phi'(\alpha)$ derivative of $\phi(\alpha)$. By using $\|\mathbf{w}\|_2 = \sum_i \frac{\gamma_i^2}{(\lambda_i + \alpha)^2}$ we can see that this derivative is

$$\phi'(\alpha) = -\frac{1}{\|\mathbf{w}\|_2^{\frac{3}{2}}} \sum_i \gamma_i^2 \frac{1}{(\lambda_i + \alpha)^2},$$

and so the update would be

$$\alpha' = \alpha + \frac{\|\mathbf{w}\|_2^2}{\sum_i \frac{\gamma_i^2}{(\lambda_i + \alpha)^3}} \cdot \frac{\|\mathbf{w}\|_2 - \lambda}{\lambda}.$$

This can be further simplified by computing the Cholesky decomposition $\mathbf{A} + \alpha\mathbf{I} = \mathbf{R}^T \mathbf{R}$, noting that $\mathbf{R}^T \mathbf{R} \mathbf{w} = -\mathbf{g}$, and defining $\mathbf{R}^T \mathbf{q} = \mathbf{w}$, as

$$\begin{aligned}
\|\mathbf{w}\|_2^2 &= \mathbf{q}^T \mathbf{R} \mathbf{R}^T \mathbf{q}, \\
&= \mathbf{q}^T \mathbf{Q}(\Lambda + \alpha\mathbf{I})\mathbf{Q}^T \mathbf{q}, \\
&= \boldsymbol{\gamma}^T (\Lambda + \alpha\mathbf{I})^{-2} \boldsymbol{\gamma}^T,
\end{aligned}$$

thus

$$\begin{aligned} \mathbf{q} &= \mathbf{Q}(\Lambda + \alpha\mathbf{I})^{-\frac{3}{2}}\boldsymbol{\gamma}, \\ \|\mathbf{q}\|_2^2 &= \boldsymbol{\gamma}^T(\Lambda + \alpha\mathbf{I})^{-\frac{3}{2}T}\mathbf{Q}^T\mathbf{Q}(\Lambda + \alpha\mathbf{I})^{-\frac{3}{2}}\boldsymbol{\gamma}, \\ &= \sum_i \frac{\gamma_i^2}{(\lambda_i + \alpha)^3}, \end{aligned}$$

and so the update is more compactly expressed as

$$\alpha' = \alpha + \frac{\|\mathbf{w}\|_2^2}{\|\mathbf{q}\|_2^2} \cdot \frac{\|\mathbf{w}\|_2 - \lambda}{\lambda}.$$

The procedure as a whole is presented in Algorithm 6. As α is initialized as 0, the first iteration will check whether the optimal w^* is in the interior of the feasible set, stopping if that is the case or moving on to an $\alpha > 0$ otherwise. As a stopping criterion we can use the distance to the boundary $|\|\mathbf{w}\|_2 - \lambda|$, except for that first iteration in which we will also stop if w is in the interior, i.e. $\|\mathbf{w}\|_2 \leq \lambda$. Note that this algorithm can be very costly for a general $\mathbf{A} + \alpha\mathbf{I}$, as Cholesky decomposition is an $O(N^3)$ procedure for the general case. Fortunately in our case the tridiagonality of \mathbf{A} allows to do this through a tridiagonal factorization (see Appendix A). In practice this method is able to converge in a very small number of iterations, thus scaling linearly with the size of the problem.

Algorithm 6 Moré-Sorensen algorithm for definite positive Hessian (optimized for tridiagonal)

Inputs: objective function in the form $f(w) = \frac{1}{2}\mathbf{w}^T\mathbf{A}\mathbf{w} + \mathbf{g}^T\mathbf{w}$, \mathbf{A} definite positive, trust region parameter λ , stopping tolerance τ .
Initialization $\alpha = 0$
while stopping criterion $> \tau$ **do**
 Compute Cholesky decomposition $\mathbf{A} + \alpha\mathbf{I} = \mathbf{R}^T\mathbf{R}$. (LAPACK DPTTRF)
 Obtain \mathbf{w} solving system $\mathbf{R}^T\mathbf{R}\mathbf{w} = -\mathbf{g}$. (LAPACK DPTTRS)
 Obtain \mathbf{q} solving system $\mathbf{R}^T\mathbf{q} = \mathbf{w}$. (LAPACK DGTSV)
 Update α as: $\alpha' = \alpha + \frac{\|\mathbf{w}\|_2^2}{\|\mathbf{q}\|_2^2} \cdot \frac{\|\mathbf{w}\|_2 - \lambda}{\lambda}$
end while

4 2-dimensional Total Variation solvers

4.1 2-dimensional ADMM

Given that the primal problem for the 2-dimensional case has a similar form to the 1-dimensional one, we can apply once again the ADMM method to solve the problem. Note however that the structure of \mathbf{B} is noticeably more complex than that of \mathbf{D} , and so some changes in the algorithm are required.

In principle, the pseudocode presented in Algorithm 2 for the 1-d case could be applied unchanged here. However, when attempting to solve the equation system induced by the matrix $\gamma\mathbf{I}_{N^2} + \mathbf{B}^T\mathbf{B}$ we can no longer rely on tridiagonality, nor precompute a tridiagonal factorization. To make up for this, note that using the Kronecker product properties we can write

$$\begin{aligned} \gamma\mathbf{I}_{N^2} + \mathbf{B}^T\mathbf{B} &= \gamma\mathbf{I}_{N^2} + [(\mathbf{I}_N \otimes \mathbf{D}^T) \quad (\mathbf{D}^T \otimes \mathbf{I}_N)] \begin{bmatrix} \mathbf{I}_N \otimes \mathbf{D} \\ \mathbf{D} \otimes \mathbf{I}_N \end{bmatrix} \\ &= \gamma\mathbf{I}_{N^2} + \mathbf{I}_N \otimes \mathbf{D}^T\mathbf{D} + \mathbf{D}^T\mathbf{D} \otimes \mathbf{I}_N \\ &= \gamma(\mathbf{I}_N \otimes \mathbf{I}_N) + \mathbf{I}_N \otimes \mathbf{D}^T\mathbf{D} + \mathbf{D}^T\mathbf{D} \otimes \mathbf{I}_N \\ &= \mathbf{I}_N \otimes (\gamma\mathbf{I}_N + \mathbf{D}^T\mathbf{D}) + \mathbf{D}^T\mathbf{D} \otimes \mathbf{I}_N \end{aligned}$$

Considering now the system we need to solve, we have $\mathbf{A}\mathbf{x} = \mathbf{c}$, where $\mathbf{A} = \gamma\mathbf{I}_{N^2} + \mathbf{B}^T\mathbf{B} = \mathbf{I}_N \otimes (\gamma\mathbf{I}_N + \mathbf{D}^T\mathbf{D}) + \mathbf{D}^T\mathbf{D} \otimes \mathbf{I}_N$ and $\mathbf{c} = \gamma\mathbf{y} + \mathbf{B}(\mathbf{p} - \mathbf{z})$. We can now make use of the Kronecker product property (2.17), obtaining

$$(\gamma\mathbf{I}_N + \mathbf{D}^T\mathbf{D})\mathbf{X} + \mathbf{X}\mathbf{D}^T\mathbf{D} = \mathbf{C}, \tag{4.1}$$

where C is c in matrix form, which can be expressed conveniently using

$$\begin{aligned}
c &= \gamma \mathbf{y} + \mathbf{B}(\mathbf{p} - \mathbf{z}) \\
&= \gamma \mathbf{y} + [(\mathbf{I}_N \otimes \mathbf{D}^T) \quad (\mathbf{D}^T \otimes \mathbf{I}_N)] (\mathbf{p} - \mathbf{z}) \\
&= \gamma \mathbf{y} + [(\mathbf{I}_N \otimes \mathbf{D}^T) \quad (\mathbf{D}^T \otimes \mathbf{I}_N)] \begin{bmatrix} \mathbf{p}_1 - \mathbf{z}_1 \\ \mathbf{p}_2 - \mathbf{z}_2 \end{bmatrix} \\
C &= \gamma \mathbf{Y} + \mathbf{D}^T(\mathbf{P}_1 - \mathbf{Z}_1) + (\mathbf{P}_2 - \mathbf{Z}_2)\mathbf{D},
\end{aligned}$$

where \mathbf{p}_1 is the vector containing the first half of \mathbf{p} , \mathbf{p}_2 the second half, equivalently for \mathbf{z}_1 , \mathbf{z}_2 , and \mathbf{P}_1 , $\mathbf{Z}_1 \in \mathcal{R}^{(N-1) \times N}$, \mathbf{P}_2 , $\mathbf{Z}_2 \in \mathcal{R}^{N \times (N-1)}$ are their respective matrix forms.

The equation system posed in (4.1) is recognised as the well-known Sylvester equation, for which a solution is known since long [2]. Given a Sylvester equation in the form $\mathbf{A}\mathbf{X} + \mathbf{X}\mathbf{B} = \mathbf{C}$ we can find the solution by first computing the Schur decomposition of matrices \mathbf{A} and \mathbf{B} as

$$\begin{aligned}
\mathbf{A} &= \mathbf{T}\mathbf{U}\mathbf{T}^T \\
\mathbf{B} &= \mathbf{Q}\mathbf{V}\mathbf{Q}^T,
\end{aligned}$$

where \mathbf{T} and \mathbf{Q} are orthogonal matrices, \mathbf{U} is lower triangular and \mathbf{V} upper triangular. A decomposition of this form is always guaranteed to exist. Now the Sylvester equation can be transformed as

$$\begin{aligned}
\mathbf{A}\mathbf{X} + \mathbf{X}\mathbf{B} &= \mathbf{C} \\
\mathbf{T}\mathbf{U}\mathbf{T}^T\mathbf{X} + \mathbf{X}\mathbf{Q}\mathbf{V}\mathbf{Q}^T &= \mathbf{C} \\
\mathbf{U}\mathbf{T}^T\mathbf{X} + \mathbf{T}^T\mathbf{X}\mathbf{Q}\mathbf{V}\mathbf{Q}^T &= \mathbf{T}^T\mathbf{C} \\
\mathbf{U}\mathbf{T}^T\mathbf{X}\mathbf{Q} + \mathbf{T}^T\mathbf{X}\mathbf{Q}\mathbf{V} &= \mathbf{T}^T\mathbf{C}\mathbf{Q}.
\end{aligned} \tag{4.2}$$

Defining now $\tilde{\mathbf{X}} = \mathbf{T}^T\mathbf{X}\mathbf{Q}$, $\tilde{\mathbf{C}} = \mathbf{T}^T\mathbf{C}\mathbf{Q}$ we have the system

$$\mathbf{U}\tilde{\mathbf{X}} + \tilde{\mathbf{X}}\mathbf{V} = \tilde{\mathbf{C}}.$$

Making use now of the triangular structure of \mathbf{U} , \mathbf{V} it is easy to see that the entries of the solution $\tilde{\mathbf{X}}$ can be computed as

$$[\tilde{\mathbf{X}}]_{ij} = \frac{[\tilde{\mathbf{C}}]_{ij} - \sum_{k=1}^{i-1} \mathbf{U}_{ik}[\tilde{\mathbf{X}}]_{kj} - \sum_{k=1}^{j-1} [\tilde{\mathbf{X}}]_{ik}\mathbf{V}_{kj}}{\mathbf{U}_{ii} + \mathbf{V}_{jj}}.$$

As for the computation of $[\tilde{\mathbf{X}}]_{ij}$ only those entries of $\tilde{\mathbf{X}}$ in previous rows/columns are needed, we can safely compute the entries of $\tilde{\mathbf{X}}$ one at a time following either a row or column order. Once $\tilde{\mathbf{X}}$ has been obtained we can recover our original problem solution \mathbf{X} by noting that by the orthogonality of \mathbf{T} and \mathbf{Q} we have $\mathbf{X} = \mathbf{T}\tilde{\mathbf{X}}\mathbf{Q}^T$.

Turning now our attention back to the equation system (4.1) we need to solve in the ADMM algorithm, the first step would be to compute the Schur decompositions of $\mathbf{D}^T\mathbf{D}$ and $\gamma\mathbf{I}_N + \mathbf{D}^T\mathbf{D}$. Both matrices are tridiagonal symmetric, and it can be shown [20] that in this case the resulting Schur decomposition is equivalent to the eigenvalue decomposition, where the eigenvectors have been made orthogonal. This fact is advantageous, as we obtain a decomposition $\mathbf{D}^T\mathbf{D} = \mathbf{T}\mathbf{U}\mathbf{T}^T$ in which \mathbf{U} is a diagonal matrix containing the eigenvalues of $\mathbf{D}^T\mathbf{D}$. What is more, the eigendecomposition of $\gamma\mathbf{I}_N + \mathbf{D}^T\mathbf{D}$ can be shown to be $\mathbf{T}(\mathbf{U} + \gamma\mathbf{I}_N)\mathbf{T}^T$. Thus, we only need to compute the eigendecomposition of the tridiagonal matrix $\mathbf{D}^T\mathbf{D}$, for which a LAPACK routine exists (DSTEV). Once that is achieved we can obtain $\tilde{\mathbf{X}}$ using (4.2), which in our particular case simplifies down to the easier formula.

$$\tilde{\mathbf{X}}_{ij} = \frac{\tilde{\mathbf{C}}_{ij}}{\mathbf{U}_{ii} + \mathbf{U}_{jj} + \gamma}, \tag{4.3}$$

hence allowing to compute $\tilde{\mathbf{X}}_{ij}$ in $O(N^2)$ time. Unfortunately, computing the Schur decomposition and afterwards performing the operation $\mathbf{X} = \mathbf{T}\tilde{\mathbf{X}}\mathbf{T}^T$ both involve a cost of order $O(N^3)$, turning this into the overall iteration

Algorithm 7 Alternating-Direction Method of Multipliers for 2-dimensional TV-L p norm

Inputs: reference data Y , augmented Lagrangian index $\gamma > 0$, stopping tolerance τ , norm order p , penalty term λ .

Initialization: set initial values for $\mathbf{P}_1, \mathbf{P}_2, \mathbf{Z}_1, \mathbf{Z}_2$.

Precompute Schur decomposition $(D^T D) = \mathbf{T} \mathbf{U} \mathbf{T}^T$. (LAPACK: DSTEVD)

while stopping criterion $> \tau$ **do**

 Compute $\mathbf{C} = D^T(\mathbf{P}_1 - \mathbf{Z}_1) + (\mathbf{P}_2 - \mathbf{Z}_2)D + \gamma Y$.

 Compute $\tilde{\mathbf{C}} = \mathbf{T}^T \mathbf{C} \mathbf{T}$.

 Solve Sylvester equation $\mathbf{U} \tilde{\mathbf{X}} + \tilde{\mathbf{X}} \mathbf{U} = \tilde{\mathbf{C}}$. (using (4.3))

 Recover $\mathbf{X} = \mathbf{T} \tilde{\mathbf{X}} \mathbf{T}^T$.

$\mathbf{S}_1 = D \mathbf{X}, \mathbf{S}_2 = \mathbf{X} D$

$\begin{bmatrix} \mathbf{P}_1 \\ \mathbf{P}_2 \end{bmatrix} = \text{prox}_{\gamma \lambda \|\cdot\|_p} \left(\begin{bmatrix} \text{vec}(\mathbf{S}_1 + \mathbf{Z}_1) \\ \text{vec}(\mathbf{S}_2 + \mathbf{Z}_2) \end{bmatrix} \right)$

$\mathbf{Z}_1 = \mathbf{Z}_1 + \mathbf{S}_1 - \mathbf{P}_1, \mathbf{Z}_2 = \mathbf{Z}_2 + \mathbf{S}_2 - \mathbf{P}_2$

end while

cost of the algorithm. While more efficient methods for operating with the eigenvectors of DD^T can be devised, it is unclear whether this can be ported to the eigenvectors of $D^T D$.

Considering all this we could already write the pseudocode of the ADMM method, replacing in Algorithm 2 the lines corresponding to the solution of the linear equation system. However, it will prove profitable to rewrite the full algorithm in terms of matrices instead of vectorized versions of them. This version is presented in Algorithm 7. Note that the original working vectors of the algorithm \mathbf{p}, \mathbf{z} and \mathbf{s} have been transformed into couples of matrices containing the first or second half of them. As it can be observed in the Algorithm, this allows to cast vector operations like $B\mathbf{x}$ into two separate, fast matrix products $\mathbf{S}_1 = D\mathbf{X}, \mathbf{S}_2 = \mathbf{X}D^T$, thus saving the need to store B explicitly in memory. Only in the proximity step it is required to transform those matrices into vector form and then reorganize the output back to matrix form.

As already stated, the $O(N^3)$ computational cost of the algorithm is dominated by the multiplications with matrix \mathbf{T} and the Schur decomposition (though this is only computed once). Regarding the number of iterations for convergence, this quantity seems to be highly sensitive to the selected value for the parameter γ , for which the optimal value depends heavily on the size of the input N and the penalty parameter λ . No easy way to choose γ is available up to our knowledge.

4.2 2-dimensional TV-L1 case

Even though Projected Newton stands among the algorithms presented for the 1-dimensional TV-L1 case, unfortunately its applicability is limited when dealing with bidimensional entries. The problem to be solved is

$$\min_{\mathbf{w}} \quad \frac{1}{2} \|\mathbf{B}^T \mathbf{w} - \mathbf{y}\|_2^2 \quad \text{s.t.} \quad -\lambda \leq \mathbf{w} \leq \lambda, \quad (4.4)$$

featuring a Hessian BB^T . In spite of this Hessian presenting a nice structure based on Kronecker products, this structure breaks when rows/columns are removed to obtain the reduced Hessian, a necessary step in the Projected Newton algorithm. Without this structure the ability to perform a Newton step in linear time is lost, thus rendering the method inefficient.

This is not the case for the SBB algorithm, which still manages to maintain a good running time per iteration. Algorithm 8 shows an outline of this method adapted for the 2-dimensional case. As expected, the only difference strives in the use of B instead of the 1-dimensional D , which can be computed efficiently in matrix form using the properties in (2.18).

4.3 2-dimensional TV-L2 case

As in the 1-dimensional case, when $p = 2$ the problem to solve is nothing but a Trust Region Subproblem, hence making Moré-Sorensen the algorithm of choice. However, in this case the problem to solve,

$$\min_{\mathbf{w}} \quad \frac{1}{2} \|\mathbf{B}^T \mathbf{w} - \mathbf{y}\|_2^2 \quad \text{s.t.} \quad \|\mathbf{w}\|_2 \leq \lambda, \quad (4.5)$$

presents a semidefinite-positive hessian BB^T , and thus it contains zero-valued eigenvalues. Recall that the Moré-Sorensen method consisted in finding the root of the function $\psi(\alpha)$, which presented second-order poles at the $-\lambda_i$

Algorithm 8 Subspace BB for 2-dimensional TV-L1

Inputs: reference data \mathbf{y} , penalty parameter λ , initial guess \mathbf{x}^0 , stopping tolerance τ .
Definitions: $([\mathbf{x}]_P)_i = \min\{\max\{\mathbf{x}_i, -\lambda\}, \lambda\}$, orthogonal projection of \mathbf{x} onto feasible set.
Initialization $\mathbf{x} = \mathbf{x}^0$.
while stopping criterion $> \tau$ **do**
 Compute gradient $\nabla f = \mathbf{B}(\mathbf{B}^T \mathbf{x} - \mathbf{y})$
 Identify set of active constraints
 $I = \{i | (-\lambda \leq \mathbf{x}_i, \nabla f(x)_i \leq 0) \text{ or } (\lambda \leq \mathbf{x}_i, \nabla f(x)_i \geq 0)\}$
 Construct subspace gradient: $[\nabla \tilde{f}]_i = [\nabla f]_i$ if $i \notin I$, 0 otherwise.
 if iteration number is odd **then**
 Compute stepsize: $\alpha = \frac{\|\nabla \tilde{f}\|_2^2}{\langle \nabla \tilde{f}, \mathbf{B}\mathbf{B}^T \nabla \tilde{f} \rangle}$.
 else
 Compute stepsize: $\alpha = \frac{\langle \nabla \tilde{f}, \mathbf{B}\mathbf{B}^T \nabla \tilde{f} \rangle}{\|\mathbf{B}\mathbf{B}^T \nabla \tilde{f}\|_2^2}$.
 end if
 $\mathbf{x} = [\mathbf{x} - \alpha \nabla f]_+$.
end while

values. Therefore one might expect $\lim_{\alpha \rightarrow 0} \psi(\alpha) = \infty$. However, we will see that in fact this does not happen in the problem at hand.

Proposition 4. *In a problem in the form (4.5), Moré and Sorensen’s “hard case” cannot happen. Furthermore $\lim_{\alpha \rightarrow 0} \psi(\alpha) = c \in \mathbb{R}$.*

Proof. The Hessian of (3.2), $\mathbf{A} = \mathbf{B}\mathbf{B}^T$, is a Gram matrix, and as such, positive-semidefinite. To see that no poles appear at zero despite the existence of $\lambda_i = 0$ eigenvalues we need to prove that $\gamma_i = 0 \forall i$ s.t. $\lambda_i = 0$. This can be achieved by using the SVD decomposition $\mathbf{B} = \mathbf{Q}\mathbf{\Sigma}\mathbf{V}^T$, where \mathbf{Q} is the eigenvector matrix of $\mathbf{B}\mathbf{B}^T = \mathbf{A}$, \mathbf{V}^T is the eigenvector matrix of $\mathbf{B}^T\mathbf{B}$, and $\mathbf{\Sigma}$ is a diagonal matrix containing the root of the eigenvalues of $\mathbf{B}\mathbf{B}^T = \mathbf{A}$ (or equivalently, $\mathbf{B}^T\mathbf{B}$). By definition, $\gamma_i = -\mathbf{y}^T \mathbf{B}^T \mathbf{Q}_i$, and applying the SVD decomposition we find

$$\begin{aligned}
 \gamma_i &= -\mathbf{y}^T \mathbf{B}^T \mathbf{Q}_i \\
 &= -\mathbf{y}^T \mathbf{V}^T \mathbf{\Sigma} \mathbf{Q}^T \mathbf{Q}_i \\
 &= -\mathbf{y}^T \mathbf{V}^T \mathbf{\Sigma} \mathbf{e}_i \\
 &= -\lambda_i (\mathbf{y}^T \mathbf{V}^T \mathbf{e}_i)
 \end{aligned}$$

where \mathbf{e}_i is an all-zeros vector except for the i -th entry, which is valued 1. From this expression it is clear that whenever $\lambda_i = 0$, $\gamma_i = 0$. Therefore there are no poles at $\alpha = 0$, and so $\lim_{\alpha \rightarrow 0} \psi(\alpha) = c \in \mathbb{R}$. Consequently, the rightmost pole is located at some point $\tilde{\alpha} < 0$, and as by the KKT conditions (3.3) $\alpha \geq 0$, then the solution $\alpha^* \in [0, \infty)$. \square

Corollary 5. *Equivalently to the 1-dimensional case, it follows that the maximum value of the $\psi(\alpha)$ function is attained at $\alpha = 0$. Therefore, for any $\lambda > \lambda_{MAX} = \psi(0)$ the solution \mathbf{w} is always the same point in the interior of the feasible set. Thus, the range of λ values generating different models is bounded by $\lambda \in [0, \lambda_{MAX}]$.*

Being again in the “easy-case” allows us to follow a similar approach to the one applied in the 1-dimensional case. In spite of this, the singularity of the Hessian \mathbf{A} produces some difficulties. To begin with, it is not possible to compute the Cholesky decomposition for a non definite-positive matrix. Furthermore, on each iteration of the algorithm we are required to solve the system $(\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)})\mathbf{w} = -\mathbf{g}$, $\mathbf{g} = \mathbf{B}\mathbf{y}$, which is singular for $\alpha = 0$. We will see how to tackle these problems.

Consider the system $(\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)})\mathbf{w} = -\mathbf{g}$ to be solved. Making use again of the properties of the Kronecker product, we find that

$$\begin{aligned}
(A + \alpha I_{2N(N-1)})\mathbf{w} &= -\mathbf{g} \\
(\mathbf{B}\mathbf{B}^T + \alpha I_{2N(N-1)})\mathbf{w} &= -\mathbf{g} \\
\left(\left[\begin{array}{c} I_N \otimes D \\ D \otimes I_N \end{array} \right] [(I_N \otimes D^T) \quad (D^T \otimes I_N)] + \alpha I_{2N(N-1)} \right) \mathbf{w} &= -\mathbf{g} \\
\left(\left[\begin{array}{cc} I_N \otimes DD^T & D^T \otimes D \\ D \otimes D^T & DD^T \otimes I_N \end{array} \right] + \alpha \left[\begin{array}{cc} I_N \otimes I_{N-1} & \\ & I_{N-1} \otimes I_N \end{array} \right] \right) \mathbf{w} &= -\mathbf{g} \\
\left[\begin{array}{cc} I_N \otimes (DD^T + \alpha I_{N-1}) & D^T \otimes D \\ D \otimes D^T & (DD^T + \alpha I_{N-1}) \otimes I_N \end{array} \right] \mathbf{w} &= -\mathbf{g} \\
\left[\begin{array}{cc} I_N \otimes (DD^T + \alpha I_{N-1}) & D^T \otimes D \\ D \otimes D^T & (DD^T + \alpha I_{N-1}) \otimes I_N \end{array} \right] \begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{bmatrix} &= - \begin{bmatrix} \mathbf{g}_1 \\ \mathbf{g}_2 \end{bmatrix},
\end{aligned}$$

where we have split \mathbf{g} and \mathbf{w} into two halves, $\mathbf{g}_1, \mathbf{g}_2$ and $\mathbf{w}_1, \mathbf{w}_2$. We can now break down the system into two smaller though related equation systems, which we can cast into matrix equation systems as follows

$$\begin{cases} (I_N \otimes (DD^T + \alpha I_{N-1}))\mathbf{w}_1 + (D^T \otimes D)\mathbf{w}_2 = -\mathbf{g}_1 \\ (D \otimes D^T)\mathbf{w}_1 + ((DD^T + \alpha I_{N-1}) \otimes I_N)\mathbf{w}_2 = -\mathbf{g}_2 \end{cases}$$

$$\begin{cases} (DD^T + \alpha I_{N-1})\mathbf{W}_1 + D\mathbf{W}_2D = -\mathbf{G}_1 \\ D^T\mathbf{W}_1D^T + \mathbf{W}_2(DD^T + \alpha I_{N-1}) = -\mathbf{G}_2 \end{cases},$$

where $\mathbf{W}_1, \mathbf{G}_1 \in \mathcal{R}^{(N-1) \times N}$, $\mathbf{W}_2, \mathbf{G}_2 \in \mathcal{R}^{N \times (N-1)}$ are the matrix versions of $\mathbf{w}_1, \mathbf{g}_1, \mathbf{w}_2, \mathbf{g}_2$. Systems of these kind are known as generalized Sylvester equations, and in the general case they lack of an easy solution [24]. Fortunately, once again the special structure of our problem makes the solving of this equation system feasible. To clarify the following derivation, we will use the definitions $M = DD^T$ and $\bar{M} = DD^T + \alpha I_{N-1}$, so that we can rewrite the equation system as

$$\begin{cases} \bar{M}\mathbf{W}_1 + D\mathbf{W}_2D = -\mathbf{G}_1 \\ D^T\mathbf{W}_1D^T + \mathbf{W}_2\bar{M} = -\mathbf{G}_2 \end{cases}$$

$$\begin{cases} \mathbf{W}_1 = \bar{M}^{-1}(-\mathbf{G}_1 - D\mathbf{W}_2D) \\ \mathbf{W}_2 = (-\mathbf{G}_2 - D^T\mathbf{W}_1D^T)\bar{M}^{-1} \end{cases},$$

introducing the second equation into the first one, we get

$$\begin{aligned}
\mathbf{W}_1 &= \bar{M}^{-1}(-\mathbf{G}_1 - D(-\mathbf{G}_2 - D^T\mathbf{W}_1D^T)\bar{M}^{-1}D) \\
\bar{M}\mathbf{W}_1 &= -\mathbf{G}_1 - D(-\mathbf{G}_2 - D^T\mathbf{W}_1D^T)\bar{M}^{-1}D \\
\bar{M}\mathbf{W}_1 &= -\mathbf{G}_1 + D\mathbf{G}_2\bar{M}^{-1}D + M\mathbf{W}_1D^T\bar{M}^{-1}D \\
\mathbf{G}_1 - D\mathbf{G}_2\bar{M}^{-1}D &= -\bar{M}\mathbf{W}_1 + M\mathbf{W}_1D^T\bar{M}^{-1}D \\
M^{-1}(\mathbf{G}_1 - D\mathbf{G}_2\bar{M}^{-1}D) &= -M^{-1}\bar{M}\mathbf{W}_1 + M^{-1}M\mathbf{W}_1D^T\bar{M}^{-1}D \\
\mathbf{C} &= -(I_{N-1} + \alpha M^{-1})\mathbf{W}_1 + \mathbf{W}_1D^T\bar{M}^{-1}D,
\end{aligned}$$

with $\mathbf{C} = M^{-1}(\mathbf{G}_1 - D\mathbf{G}_2\bar{M}^{-1}D)$, which is nothing else but an standard Sylvester equation. Suppose now that we can compute the eigendecomposition $M = \mathbf{T}\mathbf{U}\mathbf{T}^T$, and we define $\bar{\mathbf{U}} = \mathbf{U} + \alpha I_{N-1}$, so that we also have the eigendecomposition $\bar{M} = \mathbf{T}\bar{\mathbf{U}}\mathbf{T}^T$. Then the equation to solve turns out to be

$$\begin{aligned}
\mathbf{C} &= -\mathbf{T}\mathbf{U}^{-1}\bar{\mathbf{U}}\mathbf{T}^T\mathbf{W}_1 + \mathbf{W}_1D^T\mathbf{T}\bar{\mathbf{U}}^{-1}\mathbf{T}^TD \\
\mathbf{C}D^T\mathbf{T} &= -\mathbf{T}\mathbf{U}^{-1}\bar{\mathbf{U}}\mathbf{T}^T\mathbf{W}_1D^T\mathbf{T} + \mathbf{W}_1D^T\mathbf{T}\bar{\mathbf{U}}^{-1}\mathbf{T}^TDD^T\mathbf{T} \\
\mathbf{C}D^T\mathbf{T} &= -\mathbf{T}\mathbf{U}^{-1}\bar{\mathbf{U}}\mathbf{T}^T\mathbf{W}_1D^T\mathbf{T} + \mathbf{W}_1D^T\mathbf{T}\bar{\mathbf{U}}^{-1}\mathbf{U} \\
\mathbf{T}^T\mathbf{C}D^T\mathbf{T} &= -\mathbf{U}^{-1}\bar{\mathbf{U}}\mathbf{T}^T\mathbf{W}_1D^T\mathbf{T} + \mathbf{T}^T\mathbf{W}_1D^T\mathbf{T}\bar{\mathbf{U}}^{-1}\mathbf{U}.
\end{aligned}$$

Defining now

$$\begin{aligned}
\tilde{\mathbf{C}} &= \mathbf{T}^T \mathbf{C} \mathbf{D}^T \mathbf{T} \\
&= \mathbf{U}^{-1} (\mathbf{T}^T \mathbf{G}_1 - \mathbf{T}^T \mathbf{D} \mathbf{G}_2 \mathbf{T} \bar{\mathbf{U}}^{-1} \mathbf{T}^T \mathbf{D}) \mathbf{D}^T \mathbf{T} \\
&= \mathbf{U}^{-1} (\mathbf{T}^T \mathbf{G}_1 \mathbf{D}^T \mathbf{T} - \mathbf{T}^T \mathbf{D} \mathbf{G}_2 \mathbf{T} \mathbf{U} \bar{\mathbf{U}}^{-1}) \\
\tilde{\mathbf{W}}_1 &= \mathbf{T}^T \mathbf{W}_1 \mathbf{D}^T \mathbf{T},
\end{aligned} \tag{4.6}$$

we can rewrite the equation as

$$\tilde{\mathbf{C}} = -\mathbf{U}^{-1} \bar{\mathbf{U}} \tilde{\mathbf{W}}_1 + \tilde{\mathbf{W}}_1 \bar{\mathbf{U}}^{-1} \mathbf{U},$$

where the coefficient matrices multiplying $\tilde{\mathbf{W}}_1$ are diagonal. Hence we can solve for $\tilde{\mathbf{W}}_1$ in a similar fashion to our solution of the 2-dimensional ADMM method, resulting in

$$[\tilde{\mathbf{W}}_1]_{ij} = \frac{\tilde{\mathbf{C}}_{ij}}{-\mathbf{U}_{ii}^{-1} \bar{\mathbf{U}}_{ii} + \bar{\mathbf{U}}_{jj}^{-1} \mathbf{U}_{jj}} = \frac{\tilde{\mathbf{C}}_{ij}}{-1 - \frac{\alpha}{\bar{\mathbf{U}}_{ii}} + \frac{\mathbf{U}_{jj}}{\bar{\mathbf{U}}_{jj} + \alpha}}. \tag{4.7}$$

With this we can recover the original \mathbf{W}_1 , \mathbf{W}_2 solutions for our system by noting that $\mathbf{W}_1 \mathbf{D}^T = \mathbf{T} \tilde{\mathbf{W}}_1 \mathbf{T}^T$, and so

$$\mathbf{W}_2 = (-\mathbf{G}_2 - \mathbf{D} \mathbf{T} \tilde{\mathbf{W}}_1 \mathbf{T}^T) \mathbf{T} \bar{\mathbf{U}}^{-1} \mathbf{T}^T \tag{4.8}$$

$$= (-\mathbf{G}_2 \mathbf{T} - \mathbf{D} \mathbf{T} \tilde{\mathbf{W}}_1) \bar{\mathbf{U}}^{-1} \mathbf{T}^T \tag{4.9}$$

$$\mathbf{W}_1 = \mathbf{T} \bar{\mathbf{U}}^{-1} \mathbf{T}^T (-\mathbf{G}_1 - \mathbf{D} \mathbf{W}_2 \mathbf{D}). \tag{4.10}$$

Following this procedure we are able to find a solution for the Sylvester system. The most expensive part of the algorithm is computing the eigendecomposition and performing multiplications with the \mathbf{T} matrices. Fortunately these tasks can be performed on $O(N^2 \log_2 N)$ time as we show in Appendix A.

With all this in mind we can solve efficiently the equation system $(\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)}) \mathbf{w} = -\mathbf{g}$. Note however that for the Moré-Sorensen algorithm we need to compute a vector \mathbf{q} obtained by solving the system $\mathbf{R}^T \mathbf{q} = \mathbf{w}$, with \mathbf{R} Cholesky decomposition matrix of $(\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)})$. Unfortunately, as stated before, computing the Cholesky decomposition is not easy in this case, and does not even exist for $\alpha = 0$. To deal with this, note that the only purpose of \mathbf{q} in the algorithm is to compute the norm $\|\mathbf{q}\|_2^2$, and so we can write

$$\begin{aligned}
\|\mathbf{q}\|_2^2 &= (\mathbf{w}^T \mathbf{R}^{-1}) (\mathbf{R}^{-1T} \mathbf{w}) \\
&= \mathbf{w}^T (\mathbf{R}^T \mathbf{R})^{-1} \mathbf{w} \\
&= \mathbf{w}^T (\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)})^{-1} \mathbf{w} \\
&= \mathbf{w}^T \mathbf{v},
\end{aligned}$$

with $\mathbf{v} = (\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)})^{-1} \mathbf{w}$, that is, \mathbf{v} can be found by solving the system $(\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)}) \mathbf{v} = \mathbf{w}$, which has the same form as $(\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)}) \mathbf{w} = -\mathbf{g}$, and thus can be solved following the same procedure.

Algorithm 9 Moré-Sorensen algorithm for 2-dimensional TV-L2

Inputs: reference image \mathbf{Y} , trust region parameter λ , stopping tolerance τ , lower bound α_{min} .

Initialization $\alpha = \alpha_{min}$

Compute $\mathbf{g} = \mathbf{B}\mathbf{y}$.

while stopping criterion $> \tau$ **do**

 Obtain \mathbf{w} solving system $(\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)}) \mathbf{w} = -\mathbf{g}$.

 Obtain \mathbf{v} solving system $(\mathbf{A} + \alpha \mathbf{I}_{2N(N-1)}) \mathbf{v} = \mathbf{w}$.

 Update α as: $\alpha' = \alpha + \frac{\|\mathbf{w}\|_2^2}{\mathbf{w}^T \mathbf{v}} \cdot \frac{\|\mathbf{w}\|_2 - \lambda}{\lambda}$

end while

Algorithm 10 Subsystem solver for Algorithm 9

Inputs: lagrange parameter α , independent vector v .
Split independent vector v into matrices G_1, G_2 .
Compute \tilde{C} using (4.6).
Compute \tilde{W}_1 using (4.7).
Compute W_2, W_1 using (4.9,4.8).
Assemble back the solution w from W_1, W_2 .

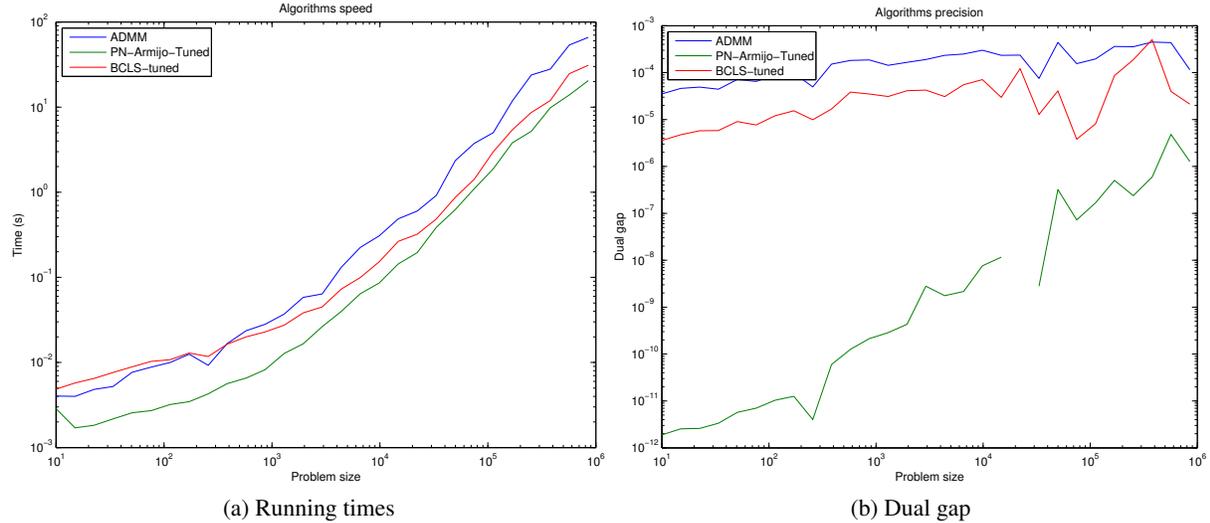


Figure 1: Running times and dual gap (accuracy) for increasing input data sizes, for the Alternating-Direction Method of Multipliers (ADMM), Projected Newton (PN-Armijo-Tuned) and SBB (BCLS-tuned), for the TV L1 1-dimensional problem

The resulting algorithm is presented as Algorithm 9, the procedure to solve the equation systems arising in each step being shown as Algorithm 10. Note that in formulas (4.6,4.7,4.8,4.9) some of the involved terms do not depend on α , and thus can be precomputed for further savings. Finally, even though the equation systems can be solved for $\alpha = 0$, this situation is prone to suffer from numerical instabilities. To overcome this we define a lower bound $\alpha_{min} \leq \alpha$ such that α is always kept in the interval $\alpha \in [\alpha_{min}, \infty)$. Even if the solution of the problem is an w generated by $\alpha = 0$, Proposition 4 guarantees that a solution for $\alpha = \alpha_{min}$ for a sufficiently small α_{min} (e.g. 10^{-7}) is a reasonable approximation, and in practice produces small enough dual gap.

5 Experiments

In this section we present some numerical results showing the performance of the developed algorithms. We will test the running times and accuracy of the solution in two different scenarios: for a fixed-magnitude penalty parameter λ and increasing sizes of the input data, and for a fixed size of the input data and increasing penalty λ . For the first scenario we will run several trials over the algorithms in which we will select a random λ in the range $[1 - 100]$, and increasing values of the input entries ranging from 10 up to 10^6 . In the second scenario we will maintain a fixed input of size 1000 and vary λ from 10^{-3} to 10^3 .

5.1 1-dimensional case

Figure 1 shows running times in seconds and dual gap obtained by the algorithms presented to solve the TV L1 problem. For ADMM a coefficient $\gamma = 1$ was selected. For Projected Newton the Armijo parameters were $\beta = 0.9$ and $\sigma = 0.1$. The three algorithms present a similar tendency in their running times as the input grows larger, though Projected Newton stands as a clear winner, both in running times and accuracy of the solution, specially in the case of small input data.

On the other hand, when the size of the input is fixed and we increase the penalty λ , large differences arise between the methods, as can be observed in Figure 2. For small penalties ADMM and specially SBB are very

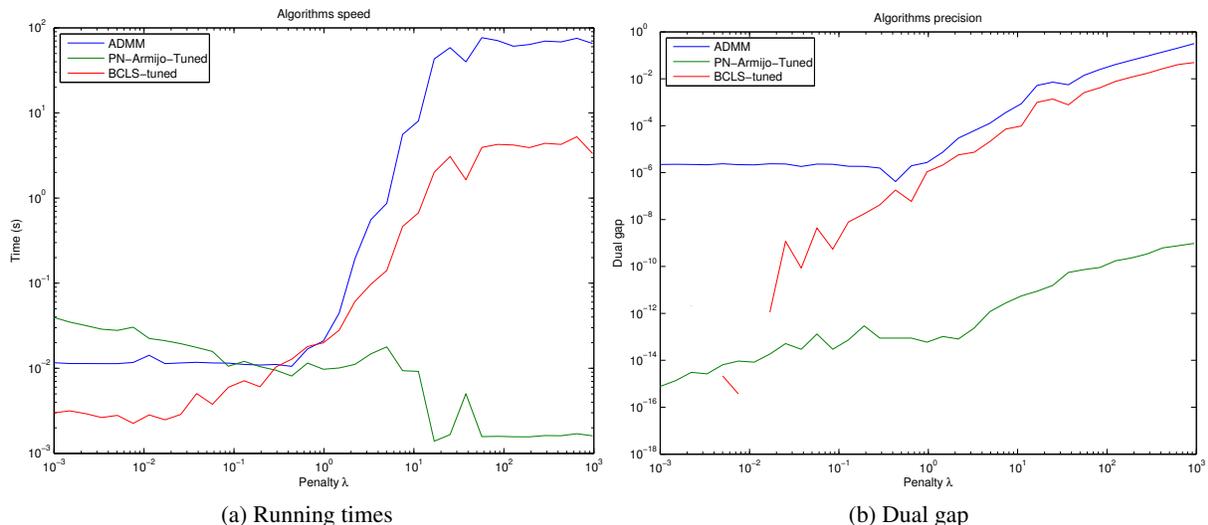


Figure 2: Running times and dual gap (accuracy) for increasing λ penalties, for the Alternating-Direction Method of Multipliers (ADMM), Projected Newton (PN-Armijo-Tuned) and SBB (BCLS-tuned), for the TV L1 1-dimensional problem.

fast, but the performance largely degrades for large penalties. Conversely Projected Newton runs fastest when the penalty is large. What is more, Projected Newton provides solutions with good accuracy for all cases. Therefore we can conclude that for this case Projected Newton is the algorithm of choice.

For L2 case we compare ADMM against Moré-Sorensen in Figure 3. For ADMM we use the rule-of-thumb of selecting $\gamma = \sqrt{N}$, which seems to work well in practice. Both methods perform adequately, though Moré-Sorensen is clearly faster and generally obtains better accuracy in spite of showing larger variance in the quality of the solutions.

Regarding the effect of λ , in Figure 4 we can clearly see that the performance of ADMM largely degrades for large values of the penalty, as well as the accuracy of the solution. Running the algorithm with different γ values somehow alleviates this problem, but there does not seem to be a clear rule for choosing γ appropriately. Conversely Moré-Sorensen seems to work well under any situation and it is parameter-free. Therefore we state that Moré-Sorensen is the best method for this case of the TV problem.

5.2 2-dimensional case

When using the L1 norm in the 2-dimensional case, the obtained results are poor. Figure 5 shows the performance of the ADMM method with parameter $\gamma = 1$. By comparing this to previous results for the 1-dimensional case, we can see that the convergence speed is slow. For the proposed extension of SBB to 2 dimensions (results not shown), even though the cost per iteration is slow, a very large number of iterations are required for convergence, turning the method impractical. In section 6 we suggest some alternative ideas to tackle this problem.

Figure 6 presents the results for Alternating-Direction Method of Multipliers (ADMM) and Moré-Sorensen algorithms in the L2 case. For ADMM a $\gamma = \sqrt{N}$ was selected. This time ADMM seems to perform better for small sizes of input data, although eventually both methods achieve similar performance. Regarding the accuracy, though, Moré-Sorensen provides better quality solutions.

It must be pointed out that most of the time required by Moré-Sorensen is spent in performing the DST operation. For this particular implementation the DST was computed by casting it into a FFT operation and using MATLAB's internal fft operation. This is not the most efficient way to implement a DFT, and thus a further speed improvement could be gained by refining this point.

When analyzing the performance of these algorithms for increasing penalty values, we observe a behaviour quite similar to the one shown for the 1-dimensional case (see Figure 7). While for small penalties both methods behave similarly, for large penalties the performance of ADMM degrades significantly. On the other hand both methods provide similar quality of solutions on average, though ADMM is more stable. Unfortunately the two of them fail to provide acceptable solutions for large penalties. Note however that such large penalties might be too strong for standard applications of the model, and so may not be as relevant in practice.

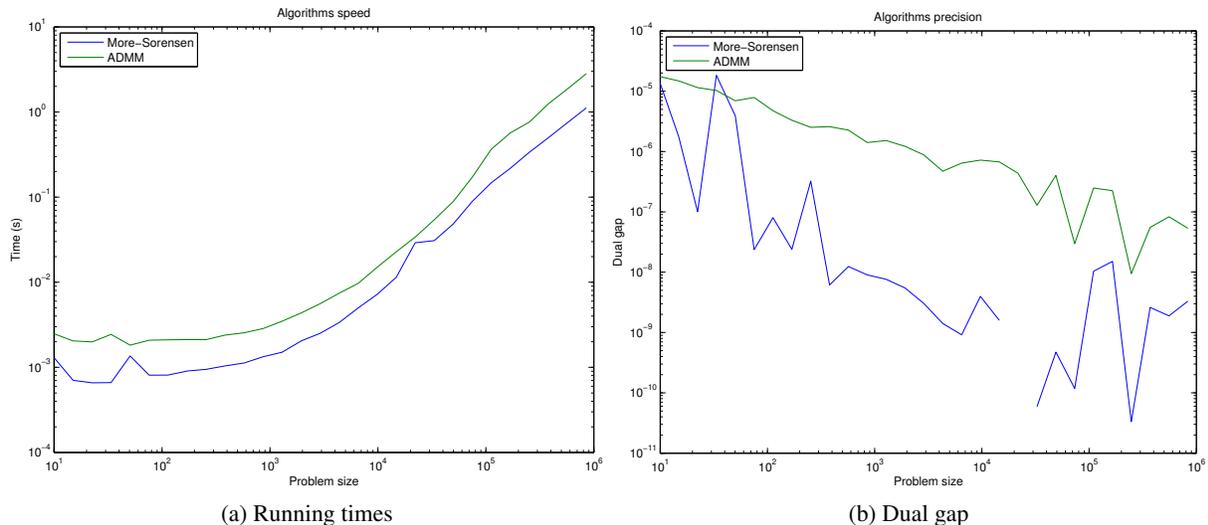


Figure 3: Running times and dual gap (accuracy) for increasing input data sizes, for the Alternating-Direction Method of Multipliers (ADMM) and Moré-Sorensen, for the TV L2 1-dimensional problem.

6 Extensions and further work

More complex models making use of a Total Variation regularization could be solved by using the proposed methods as intermediate solvers. For instance, the function

$$\min_{\mathbf{x}} f(\mathbf{x}) + \lambda \|\mathbf{M}\mathbf{x}\|_p,$$

either for $\mathbf{M} = \mathbf{D}$ or $\mathbf{M} = \mathbf{B}$ could be solved by a Trust Region algorithm like the one presented in [16], where at each step a quadratic model approximating f is constructed and minimized by solving a proximity operator in the form

$$\arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{M}\mathbf{x}\|_p,$$

for a certain \mathbf{y} , which is exactly the problem we have solved here. In this kind of setting the ability to solve fastly each of these subproblems is crucial for the performance of the overall algorithm. Having proven that our proposed algorithms are able to solve this class of problems efficiently and to a good degree of accuracy, we believe that a Trust Region approach based on them could provide good results also for this general problems.

Other proposals for TV-like problems are found in the literature, e.g. in the context of image deblurring. In [7] an L1 norm is proposed for the fidelity term instead of the standard L2 norm, resulting in an optimization problem in the form

$$\min_{\mathbf{x}} \|\mathbf{x} - \mathbf{y}\|_1 + \lambda \|\mathbf{D}\mathbf{x}\|_1.$$

Another proposal in [15] uses the p power of the TV L_p norm, resulting in

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{D}\mathbf{x}\|_p^p.$$

In [5] yet more complex extensions of the TV operator are presented. Whether the algorithms proposed in this work can be extended to these setting is an open question at the moment.

Additionally it would be of interest to test whether an Interior Point (IP) method [25] could overperform the presented algorithms for the case of the L1 norm. These algorithms are known for their good convergence properties and their flexibility to tackle a broad range of optimization problems. However, their convergence speed largely depends on the ability to solve efficiently a relaxed KKT satisfiability equation system. Software libraries such as CVXOPT [9] allow to plug-in customized solvers for this system, so that the IP method can take advantage of it. A

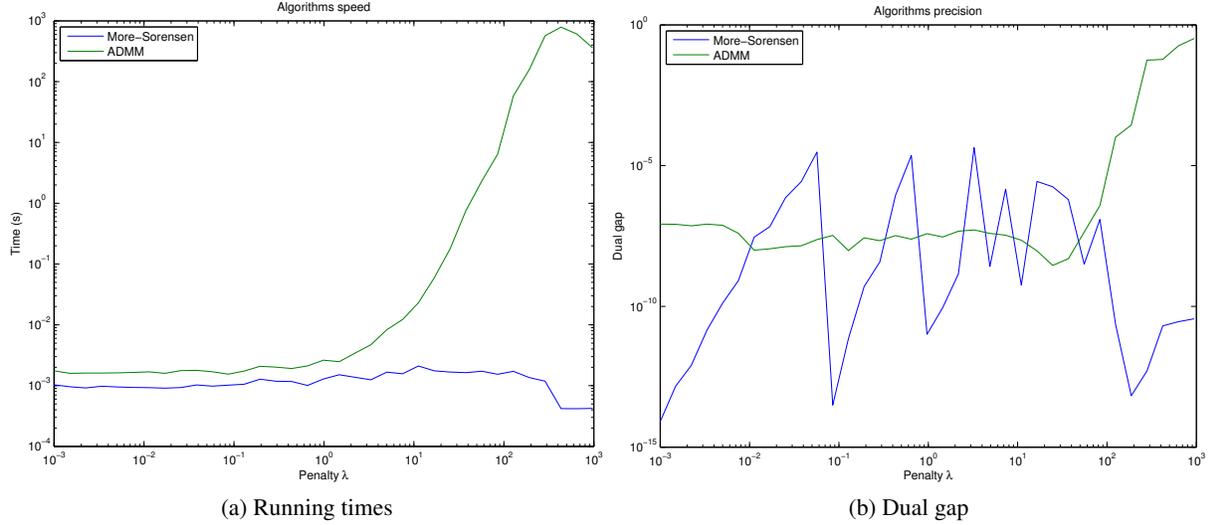


Figure 4: Running times and dual gap (accuracy) for increasing λ penalties, for the Alternating-Direction Method of Multipliers (ADMM) and Moré-Sorensen, for the TV L2 1-dimensional problem.

preliminary study has revealed that a TV L1 solver based on this approach could be feasible, though further work is needed to construct a functional implementation. Furthermore whether this approach could be extended to the L1 2-dimensional case, where we lack of an efficient algorithm, is another point of interest.

Finally, as already mentioned, the most frequent application of the TV regularization is within the context of image processing, in which it is used for operations such as denoising, inpainting and deblurring [10]. For these applications a number of software tools are already publicly available. Therefore, to correctly assess the performance of the algorithms presented here a comparison against those implementations should be made.

A Fast matrix factorizations

Matrices DD^T and $D^T D$ present the following tridiagonal structure

$$DD^T = \begin{pmatrix} 2 & -1 & & & & \\ -1 & 2 & -1 & & & \\ & -1 & 2 & -1 & & \\ & & & \ddots & & \\ & & & -1 & 2 & -1 \\ & & & & -1 & 2 \end{pmatrix} \in \mathbb{R}^{(N-1) \times (N-1)},$$

$$D^T D = \begin{pmatrix} 1 & -1 & & & & \\ -1 & 2 & -1 & & & \\ & -1 & 2 & -1 & & \\ & & & \ddots & & \\ & & & -1 & 2 & -1 \\ & & & & -1 & 1 \end{pmatrix} \in \mathbb{R}^{N \times N}.$$

Any tridiagonal matrix can be factorized in the form $M = L\Delta L^T$ [14], the factors having the following structure

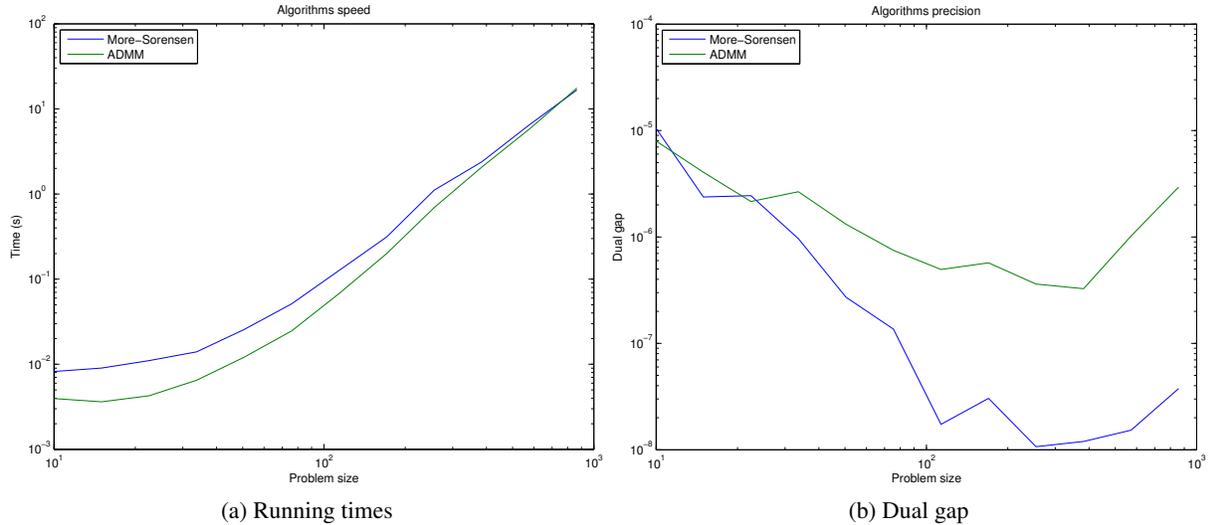


Figure 6: Running times and dual gap (accuracy) for increasing input data sizes ($N \times N$), for the Alternating-Direction Method of Multipliers (ADMM) and Moré-Sorensen, for the TV L2 2-dimensional problem.

$$[DST(\mathbf{x})]_i = \sum_{j=1}^n \mathbf{x}_j \sin\left(\frac{ij\pi}{n+1}\right).$$

By using the transformation $\tilde{\mathbf{v}}_i = \mathbf{v}_i \beta_i$, which can be written as $\tilde{\mathbf{v}} = \text{diag}(\beta)\mathbf{v}$ we find out that

$$[\mathbf{T}\mathbf{v}]_i = [DST(\tilde{\mathbf{v}})]_i = [DST(\text{diag}(\beta)\mathbf{v})]_i.$$

Therefore we can compute each resulting row of the product by matrix \mathbf{T} by performing a DST instead. The DST operation presents symmetries that allow the use of a divide-and-conquer strategy to speed-up its computation, much in the way the Fast Fourier Transform is performed, resulting in computational costs of order $O(N \log_2 N)$. Publicly available libraries like FFTW [12] offer efficient methods to perform this computation. All other possible products involving \mathbf{T} can be computed similarly, namely

$$\begin{aligned} [v^T \mathbf{T}]_i &= [\text{diag}(\beta) DST(\mathbf{v})]_i^T, \\ [\mathbf{T}^T v]_i &= [\text{diag}(\beta) DST(\mathbf{v})]_i, \\ [v^T \mathbf{T}^T]_i &= [DST(\text{diag}(\beta)\mathbf{v})]_i^T. \end{aligned}$$

Therefore a computation in the form $\tilde{m}\mathbf{v}$ results in a cost of order $O(N^2 \log_2(N))$. Note also that the Schur decomposition of a symmetric matrix (like $\mathbf{D}\mathbf{D}^T$) coincides with its eigendecomposition, and thus can be computed in the same way.

For the case of the matrix $\mathbf{D}^T \mathbf{D}$, it can be shown that the eigenvalues are the same as those for $\mathbf{D}\mathbf{D}^T$ plus a zero eigenvalue. Unfortunately there is no simple rule to compute the eigenvectors (up to our knowledge), and thus if the eigenvectors are required a general eigendecomposition for tridiagonal matrices (available in LAPACK) must be performed.

References

- [1] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999. ISBN 0-89871-447-8 (paperback). 22
- [2] R. H. Bartels and G. W. Stewart. Solution of the matrix equation $\mathbf{A}\mathbf{X} + \mathbf{X}\mathbf{B} = \mathbf{C}$. *Communications of the ACM*, 15(9):820–826, September 1972. 13

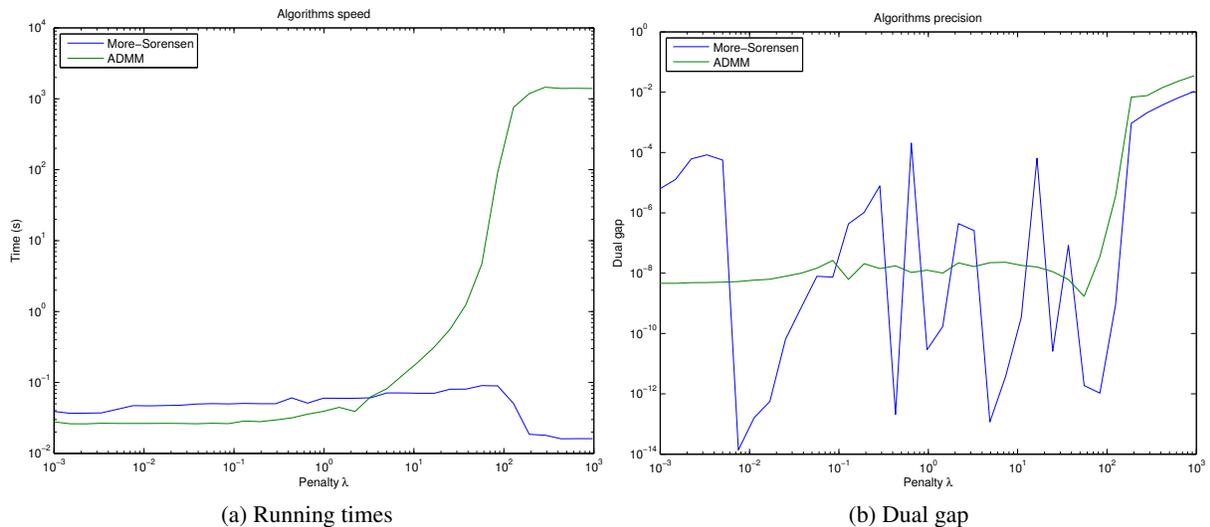


Figure 7: Running times and dual gap (accuracy) for increasing λ penalties, for the Alternating-Direction Method of Multipliers (ADMM) and Moré-Sorensen, for the TV L2 2-dimensional problem.

- [3] Jonathan Barzilai and Jonathan M. Borwein. Two-point step size gradient methods. *IMA Journal of Numerical Analysis*, 8:141–148, 1988. 9
- [4] Dimitri P. Bertsekas. Projected newton methods for optimization problems with simple constraints. *SIAM J. Control and Optimization*, 20(2), March 1982. 6, 7
- [5] Antoni Buades, Triet M. Le, Jean-Michel Morel, and Luminita A. Vese. Fast cartoon + texture image filters. *IEEE Transactions on Image Processing*, 19(8), August 2010. 20
- [6] Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. Technical report, Northwestern University, 1994. 6
- [7] Tony F. Chan and Selim Esedoglu. Aspects of total variation regularized l1 function approximation. Technical report, UCLA Mathematics Department, February 2004. 20
- [8] Patrick L. Combettes and Jean-Christophe Pesquet. Proximal splitting methods in signal processing. *arXiv*, 2009. 3, 5, 6
- [9] Joachim Dahl and Lieven Vandenbergh. Cvxopt: Python software for convex optimization. URL <http://abel.ee.ucla.edu/cvxopt/index.html>. 20
- [10] Joachim Dahl, Per Christian Hansen, Søren Holdt Jensen, and Tobias Lindstrøm Jensen. Algorithms and software for total variation image reconstruction via first-order methods. *Numer Algor*, (53):67–92, 2010. 21
- [11] Jennifer B. Erway and Philip E. Gill. A subspace minimization method for the trust-region step. *SIAM J. Optim.*, 20(3):1439–1461, 2009. 10
- [12] Matteo Frigo and Steven G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. Special issue on “Program Generation, Optimization, and Platform Adaptation”. 23
- [13] Donald Goldfarb and Wotao Yin. Second-order cone programming methods for total variation-based image restoration. *SIAM J. Sci. Comput.*, 27(2):622–645, 2005. 1
- [14] Gene H. Golub and Gerard Meurant. *Matrices, Moments and Quadrature with Applications*. Princeton University Press, 2010. 21
- [15] Christian Hansen, James G. Nagy, and Dianne P. O’Leary. *Deblurring images: matrices, spectra and filtering*. SIAM, 2006. 20
- [16] Dongmin Kim, Suvrit Sra, and Inderjit Dhillon. A scalable trust-region algorithm with application to mixed-norm regression. In *Proceedings of the 27th International Conference on Machine Learning*, 2010. 3, 20

- [17] Dongmin Kim, Suvrit Sra, and Inderjit S. Dhillon. A non-monotonic method for large-scale nonnegative least squares. Technical report, University of Texas at Austin and Max-Planck-Institute for Biological Cybernetics, May 2010. [6](#), [9](#)
- [18] Chih-Jen Lin and Jorge J. Moré. Newton’s method for large bound-constrained optimization problems. *SIAM J. Optim.*, 9(4):1100–1127, 1999. [6](#)
- [19] Jorge J. Moré and D. C. Sorensen. Computing a trust region step. *SIAM J. Sci. Stat. Comput.*, 4(3), September 1983. [10](#), [11](#)
- [20] Dianne P. O’Leary. *Scientific computing with case studies*. SIAM, 2009. [13](#), [22](#)
- [21] Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1997. [3](#)
- [22] Marielba Rojas, Sandra A. Santos, and Danny C. Sorensen. Algorithm 873: LSTRS: MATLAB software for large-scale trust-region subproblems and regularization. *ACM Trans. Math. Software*, 34(2):11, 2008. [10](#)
- [23] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992. [1](#)
- [24] Charles F. Van Loan. The ubiquitous kronecker product. *Journal of Computational and Applied Mathematics*, 123:85–100, 2000. [5](#), [16](#)
- [25] Stephen J. Wright. *Primal-Dual Interior Point Methods*. SIAM, 1997. [20](#)