
Network Flow Algorithms for Structured Sparsity

Julien Mairal*

INRIA - Willow Project-Team[†]
julien.mairal@inria.fr

Rodolphe Jenatton*

INRIA - Willow Project-Team[†]
rodolphe.jenatton@inria.fr

Guillaume Obozinski

INRIA - Willow Project-Team[†]
guillaume.obozinski@inria.fr

Francis Bach

INRIA - Willow Project-Team[†]
francis.bach@inria.fr

Abstract

We consider a class of learning problems that involve a structured sparsity-inducing norm defined as the sum of ℓ_∞ -norms over groups of variables. Whereas a lot of effort has been put in developing fast optimization methods when the groups are disjoint or embedded in a specific hierarchical structure, we address here the case of general overlapping groups. To this end, we show that the corresponding optimization problem is related to network flow optimization. More precisely, the proximal problem associated with the norm we consider is dual to a *quadratic min-cost flow problem*. We propose an efficient procedure which computes its solution exactly in polynomial time. Our algorithm scales up to millions of variables, and opens up a whole new range of applications for structured sparse models. We present several experiments on image and video data, demonstrating the applicability and scalability of our approach for various problems.

1 Introduction

Sparse linear models have become a popular framework for dealing with various unsupervised and supervised tasks in machine learning and signal processing. In such models, linear combinations of small sets of variables are selected to describe the data. Regularization by the ℓ_1 -norm has emerged as a powerful tool for addressing this combinatorial variable selection problem, relying on both a well-developed theory (see [1] and references therein) and efficient algorithms [2, 3, 4].

The ℓ_1 -norm primarily encourages sparse solutions, regardless of the potential structural relationships (e.g., spatial, temporal or hierarchical) existing between the variables. Much effort has recently been devoted to designing sparsity-inducing regularizations capable of encoding higher-order information about allowed patterns of non-zero coefficients [5, 6, 7, 8, 9, 10], with successful applications in bioinformatics [6, 11], topic modeling [12] and computer vision [9, 10].

By considering sums of norms of appropriate subsets, or *groups*, of variables, these regularizations control the sparsity patterns of the solutions. The underlying optimization problem is usually difficult, in part because it involves nonsmooth components. Proximal methods have proven to be effective in this context, essentially because of their fast convergence rates and their scalability [3, 4]. While the settings where the penalized groups of variables do not overlap or are embedded in a tree-shaped hierarchy [12] have already been studied, regularizations with general overlapping groups have, to the best of our knowledge, never been addressed with proximal methods.

This paper makes the following contributions:

- It shows that the *proximal operator* associated with the structured norm we consider can be

*Contributed equally.

[†]Laboratoire d'Informatique de l'École Normale Supérieure (INRIA/ENS/CNRS UMR 8548)

computed with a fast and scalable procedure by solving a *quadratic min-cost flow* problem.

- It shows that the dual norm of the sparsity-inducing norm we consider can also be evaluated efficiently, which enables us to compute duality gaps for the corresponding optimization problems.

- It demonstrates that our method is relevant for various applications, from video background subtraction to estimation of hierarchical structures for dictionary learning of natural image patches.

2 Structured Sparse Models

We consider in this paper convex optimization problems of the form

$$\min_{\mathbf{w} \in \mathbb{R}^p} f(\mathbf{w}) + \lambda \Omega(\mathbf{w}), \quad (1)$$

where $f : \mathbb{R}^p \rightarrow \mathbb{R}$ is a convex differentiable function and $\Omega : \mathbb{R}^p \rightarrow \mathbb{R}$ is a convex, nonsmooth, sparsity-inducing regularization function. When one knows *a priori* that the solutions of this learning problem have only a few non-zero coefficients, Ω is often chosen to be the ℓ_1 -norm (see [1, 2]). When these coefficients are organized in groups, a penalty encoding explicitly this prior knowledge can improve the prediction performance and/or interpretability of the learned models [13, 14]. Denoting by \mathcal{G} a set of groups of indices, such a penalty might for example take the form:

$$\Omega(\mathbf{w}) \triangleq \sum_{g \in \mathcal{G}} \eta_g \max_{j \in g} |\mathbf{w}_j| = \sum_{g \in \mathcal{G}} \eta_g \|\mathbf{w}_g\|_\infty, \quad (2)$$

where \mathbf{w}_j is the j -th entry of \mathbf{w} for j in $[1; p] \triangleq \{1, \dots, p\}$, the vector \mathbf{w}_g in $\mathbb{R}^{|g|}$ records the coefficients of \mathbf{w} indexed by g in \mathcal{G} , and the scalars η_g are positive weights. A sum of ℓ_2 -norms is also used in the literature [7], but the ℓ_∞ -norm is piecewise linear, a property that we take advantage of in this paper. Note that when \mathcal{G} is the set of singletons of $[1; p]$, we get back the ℓ_1 -norm.

If \mathcal{G} is a more general *partition* of $[1; p]$, variables are selected in groups rather than individually. When the groups overlap, Ω is still a norm and sets groups of variables to zero together [5]. The latter setting has first been considered for hierarchies [7, 11, 15], and then extended to general group structures [5].¹ Solving Eq. (1) in this context becomes challenging and is the topic of this paper. Following Jenatton et al. [12] who tackled the case of hierarchical groups, we propose to approach this problem with proximal methods, which we now introduce.

2.1 Proximal Methods

In a nutshell, proximal methods can be seen as a natural extension of gradient-based techniques, and they are well suited to minimizing the sum $f + \lambda \Omega$ of two convex terms, a smooth function f —continuously differentiable with Lipschitz-continuous gradient— and a potentially non-smooth function $\lambda \Omega$ (see [16] and references therein). At each iteration, the function f is linearized at the current estimate \mathbf{w}_0 and the so-called *proximal* problem has to be solved:

$$\min_{\mathbf{w} \in \mathbb{R}^p} f(\mathbf{w}_0) + (\mathbf{w} - \mathbf{w}_0)^\top \nabla f(\mathbf{w}_0) + \lambda \Omega(\mathbf{w}) + \frac{L}{2} \|\mathbf{w} - \mathbf{w}_0\|_2^2.$$

The quadratic term keeps the solution in a neighborhood where the current linear approximation holds, and $L > 0$ is an upper bound on the Lipschitz constant of ∇f . This problem can be rewritten as

$$\min_{\mathbf{w} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{w}\|_2^2 + \lambda' \Omega(\mathbf{w}), \quad (3)$$

with $\lambda' \triangleq \lambda/L$, and $\mathbf{u} \triangleq \mathbf{w}_0 - \frac{1}{L} \nabla f(\mathbf{w}_0)$. We call *proximal operator* associated with the regularization $\lambda' \Omega$ the function that maps a vector \mathbf{u} in \mathbb{R}^p onto the (unique, by strong convexity) solution \mathbf{w}^* of Eq. (3). Simple proximal methods use \mathbf{w}^* as the next iterate, but accelerated variants [3, 4] are also based on the proximal operator and require to solve problem (3) *exactly* and *efficiently* to enjoy their fast convergence rates. Note that when Ω is the ℓ_1 -norm, the solution of Eq. (3) is obtained by soft-thresholding [16]. The approach we develop in the rest of this paper extends [12] to the case of general overlapping groups when Ω is a weighted sum of ℓ_∞ -norms, broadening the application of these regularizations to a wider spectrum of problems.²

¹Note that other types of structured sparse models have also been introduced, either through a different norm [6], or through non-convex criteria [8, 9, 10].

²For hierarchies, the approach of [12] applies also to the case of where Ω is a weighted sum of ℓ_2 -norms.

3 A Quadratic Min-Cost Flow Formulation

In this section, we show that a convex dual of problem (3) for general overlapping groups \mathcal{G} can be reformulated as a *quadratic min-cost flow problem*. We present an efficient algorithm to solve it *exactly*, as well as a related algorithm to compute the dual norm of Ω . We start by considering the dual formulation to problem (3) introduced in [12], for the case where Ω is a sum of ℓ_∞ -norms:

Lemma 1 (Dual of the proximal problem [12])

Given \mathbf{u} in \mathbb{R}^p , consider the problem

$$\min_{\xi \in \mathbb{R}^{p \times |\mathcal{G}|}} \frac{1}{2} \left\| \mathbf{u} - \sum_{g \in \mathcal{G}} \xi^g \right\|_2^2 \quad \text{s.t.} \quad \forall g \in \mathcal{G}, \|\xi^g\|_1 \leq \lambda \eta_g \quad \text{and} \quad \xi_j^g = 0 \text{ if } j \notin g, \quad (4)$$

where $\xi = (\xi^g)_{g \in \mathcal{G}}$ is in $\mathbb{R}^{p \times |\mathcal{G}|}$, and ξ_j^g denotes the j -th coordinate of the vector ξ^g . Then, every solution $\xi^* = (\xi^{*g})_{g \in \mathcal{G}}$ of Eq. (4) satisfies $\mathbf{w}^* = \mathbf{u} - \sum_{g \in \mathcal{G}} \xi^{*g}$, where \mathbf{w}^* is the solution of Eq. (3).

Without loss of generality,³ we assume from now on that the scalars \mathbf{u}_j are all non-negative, and we constrain the entries of ξ to be non-negative. We now introduce a graph modeling of problem (4).

3.1 Graph Model

Let G be a directed graph $G = (V, E, s, t)$, where V is a set of vertices, $E \subseteq V \times V$ a set of arcs, s a source, and t a sink. Let c and c' be two functions on the arcs, $c : E \rightarrow \mathbb{R}$ and $c' : E \rightarrow \mathbb{R}^+$, where c is a *cost function* and c' is a non-negative *capacity function*. A *flow* is a non-negative function on arcs that satisfies capacity constraints on all arcs (the value of the flow on an arc is less than or equal to the arc capacity) and conservation constraints on all vertices (the sum of incoming flows at a vertex is equal to the sum of outgoing flows) except for the source and the sink.

We introduce a *canonical* graph G associated with our optimization problem, and uniquely characterized by the following construction:

- (i) V is the union of two sets of vertices V_u and V_{gr} , where V_u contains exactly one vertex for each index j in $[1; p]$, and V_{gr} contains exactly one vertex for each group g in \mathcal{G} . We thus have $|V| = |\mathcal{G}| + p$. For simplicity, we identify groups and indices with the vertices of the graph.
- (ii) For every group g in \mathcal{G} , E contains an arc (s, g) . These arcs have capacity $\lambda \eta_g$ and zero cost.
- (iii) For every group g in \mathcal{G} , and every index j in g , E contains an arc (g, j) with zero cost and infinite capacity. We denote by ξ_j^g the flow on this arc.
- (iv) For every index j in $[1; p]$, E contains an arc (j, t) with infinite capacity and a cost $c_j \triangleq \frac{1}{2}(\mathbf{u}_j - \bar{\xi}_j)^2$, where $\bar{\xi}_j$ is the flow on (j, t) . Note that by flow conservation, we necessarily have $\bar{\xi}_j = \sum_{g \in \mathcal{G}} \xi_j^g$.

Examples of canonical graphs are given in Figures 1(a)-(c). The flows ξ_j^g associated with G can now be identified with the variables of problem (4): indeed, the sum of the costs on the edges leading to the sink is equal to the objective function of (4), while the capacities of the arcs (s, g) match the constraints on each group. This shows that finding a flow *minimizing the sum of the costs* on such a graph is equivalent to solving problem (4).

When some groups are included in others, the canonical graph can be simplified to yield a graph with a smaller number of edges. Specifically, if h and g are groups with $h \subset g$, the edges (g, j) for $j \in h$ carrying a flow ξ_j^g can be removed and replaced by a single edge (g, h) of infinite capacity and zero cost, carrying the flow $\sum_{j \in h} \xi_j^g$. This simplification is illustrated in Figure 1(d), with a graph equivalent to the one of Figure 1(c). This does not change the optimal value of $\bar{\xi}^*$, which is the quantity of interest for computing the optimal primal variable \mathbf{w}^* (a proof and a formal definition of these equivalent graphs are available in a longer technical report [17]). These simplifications are useful in practice, since they reduce the number of edges in the graph and improve the speed of the algorithms we are now going to present.

³Let ξ^* denote a solution of Eq. (4). Optimality conditions of Eq. (4) derived in [12] show that for all j in $[1; p]$, the signs of the non-zero coefficients ξ_j^{*g} for g in \mathcal{G} are the same as the signs of the entries \mathbf{u}_j . To solve Eq. (4), one can therefore flip the signs of the negative variables \mathbf{u}_j , then solve the modified dual formulation (with non-negative variables), which gives the magnitude of the entries ξ_j^{*g} (the signs of these being known).

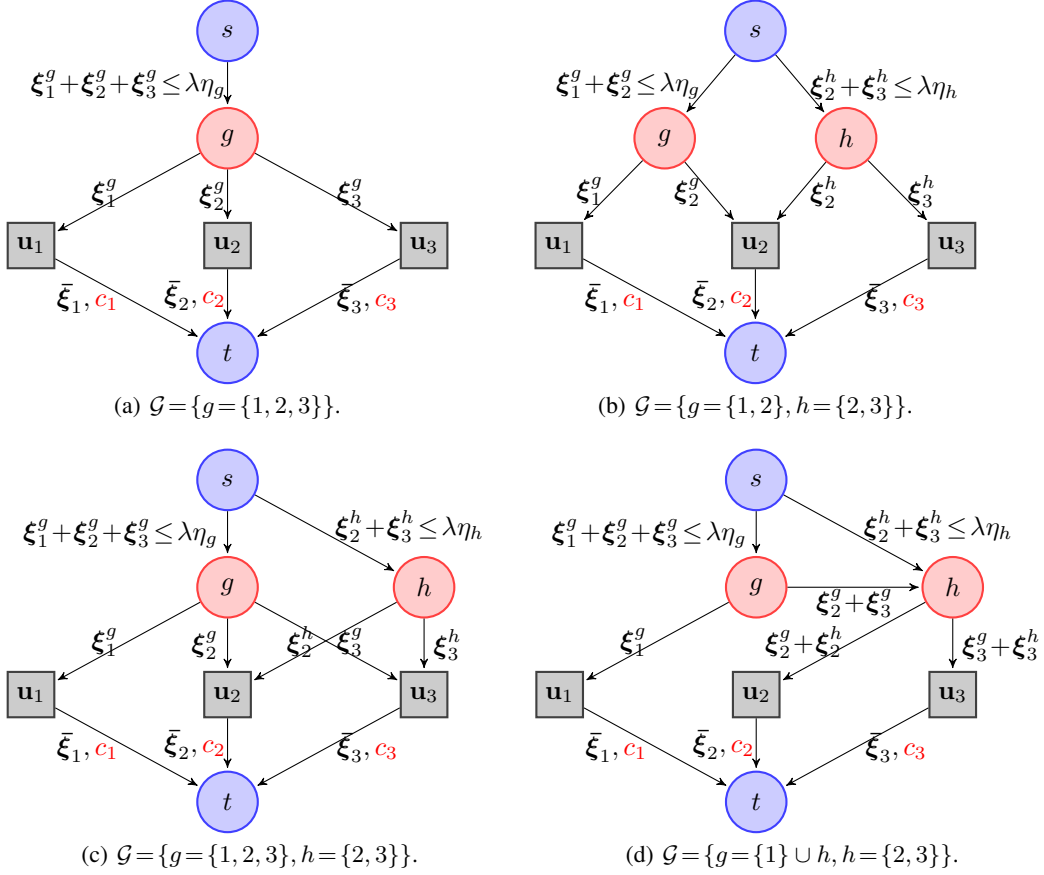


Figure 1: Graph representation of simple proximal problems with different group structures \mathcal{G} . The three indices 1, 2, 3 are represented as grey squares, and the groups g, h in \mathcal{G} as red discs. The source is linked to every group g, h with respective maximum capacity $\lambda\eta_g, \lambda\eta_h$ and zero cost. Each variable \mathbf{u}_j is linked to the sink t , with an infinite capacity, and with a cost $c_j \triangleq \frac{1}{2}(\mathbf{u}_j - \bar{\xi}_j)^2$. All other arcs in the graph have zero cost and infinite capacity. They represent inclusion relationships in-between groups, and between groups and variables. The graphs (c) and (d) correspond to a special case of tree-structured hierarchy in the sense of [12]. Their min-cost flow problems are equivalent.

3.2 Computation of the Proximal Operator

Quadratic min-cost flow problems have been well studied in the operations research literature [18]. One of the simplest cases, where \mathcal{G} contains a single group g (Ω is the ℓ_∞ -norm) as in Figure 1(a), can be solved by an orthogonal projection on the ℓ_1 -ball of radius $\lambda\eta_g$. It has been shown that such a projection can be done in $O(p)$ operations [18, 19]. When the group structure is a tree as in Figure 1(d), the problem can be solved in $O(pd)$ operations, where d is the depth of the tree [12, 18].⁴

The general case of overlapping groups is more difficult. Hochbaum and Hong have shown in [18] that *quadratic min-cost flow problems* can be reduced to a specific *parametric max-flow problem*, for which an efficient algorithm exists [20].⁵ While this generic approach could be used to solve Eq. (4), we propose to use Algorithm 1 that also exploits the fact that our graphs have non-zero costs only on edges leading to the sink. As shown in the technical report [17], it has a significantly better performance in practice. This algorithm clearly shares some similarities with existing approaches in network flow optimization such as the simplified version of [20] presented in [21] that uses a divide and conquer strategy. Moreover, we have discovered after that this paper was accepted for publication that an equivalent algorithm exists for minimizing convex functions over polymatroid

⁴When restricted to the case where Ω is a sum of ℓ_∞ -norms, the approach of [12] is in fact similar to [18].

⁵By definition, a parametric max-flow problem consists in solving, for every value of a parameter, a max-flow problem on a graph whose arc capacities depend on this parameter.

sets [22]. This equivalence, however, requires a non-trivial representation of structured sparsity-inducing norms with submodular functions, as recently pointed out by [23].

Algorithm 1 Computation of the proximal operator for overlapping groups.

- 1: **Inputs:** $\mathbf{u} \in \mathbb{R}^p$, a set of groups \mathcal{G} , positive weights $(\eta_g)_{g \in \mathcal{G}}$, and λ (regularization parameter).
- 2: Build the initial graph $G_0 = (V_0, E_0, s, t)$ as explained in Section 3.2.
- 3: Compute the optimal flow: $\bar{\xi} \leftarrow \text{computeFlow}(V_0, E_0)$.
- 4: **Return:** $\mathbf{w} = \mathbf{u} - \bar{\xi}$ (optimal solution of the proximal problem).

Function `computeFlow`($V = V_u \cup V_{gr}, E$)

- 1: Projection step: $\gamma \leftarrow \arg \min_{\gamma} \sum_{j \in V_u} \frac{1}{2}(\mathbf{u}_j - \gamma_j)^2$ s.t. $\sum_{j \in V_u} \gamma_j \leq \lambda \sum_{g \in V_{gr}} \eta_g$.
 - 2: For all nodes j in V_u , set γ_j to be the capacity of the arc (j, t) .
 - 3: Max-flow step: Update $(\bar{\xi}_j)_{j \in V_u}$ by computing a max-flow on the graph (V, E, s, t) .
 - 4: **if** $\exists j \in V_u$ s.t. $\bar{\xi}_j \neq \gamma_j$ **then**
 - 5: Denote by (s, V^+) and (V^-, t) the two disjoint subsets of (V, s, t) separated by the minimum (s, t) -cut of the graph, and remove the arcs between V^+ and V^- . Call E^+ and E^- the two remaining disjoint subsets of E corresponding to V^+ and V^- .
 - 6: $(\bar{\xi}_j)_{j \in V_u^+} \leftarrow \text{computeFlow}(V^+, E^+)$.
 - 7: $(\bar{\xi}_j)_{j \in V_u^-} \leftarrow \text{computeFlow}(V^-, E^-)$.
 - 8: **end if**
 - 9: **Return:** $(\bar{\xi}_j)_{j \in V_u}$.
-

The intuition behind this algorithm is the following: The first step looks for a candidate value for $\bar{\xi} = \sum_{g \in \mathcal{G}} \xi^g$ by solving a relaxed version of problem Eq. (4), where the constraints $\|\xi^g\|_1 \leq \lambda \eta_g$ are dropped and replaced by a single one $\|\bar{\xi}\|_1 \leq \lambda \sum_{g \in \mathcal{G}} \eta_g$. The relaxed problem only depends on $\bar{\xi}$ and can be solved in linear time. By calling its solution γ , it provides a lower bound $\|\mathbf{u} - \gamma\|_2^2/2$ on the optimal cost. Then, the second step tries to find a feasible flow of the original problem (4) such that the resulting vector $\bar{\xi}$ matches γ , which is in fact a max-flow problem [24]. If $\bar{\xi} = \gamma$, then the cost of the flow reaches the lower bound, and the flow is optimal. If $\bar{\xi} \neq \gamma$, the lower bound is not achievable, and we construct a minimum (s, t) -cut of the graph [25] that defines two disjoint sets of nodes V^+ and V^- ; V^+ is the part of the graph that could potentially have received more flow from the source (the arcs between s and V^+ are not saturated), whereas all arcs linking s to V^- are saturated. At this point, it is possible to show that the value of the optimal min-cost flow on all arcs between V^+ and V^- is necessary zero. Thus, removing them yields an equivalent optimization problem, which can be decomposed into two independent problems of smaller sizes and solved recursively by the calls to `computeFlow`(V^+, E^+) and `computeFlow`(V^-, E^-). A formal proof of correctness of Algorithm 1 and further details are relegated to [17].

The approach of [18, 20] is guaranteed to have the same worst-case complexity as a single max-flow algorithm. However, we have experimentally observed a significant discrepancy between the worst case and empirical complexities for these flow problems, essentially because the empirical cost of each max-flow is significantly smaller than its theoretical cost. Despite the fact that the worst-case guarantee of our algorithm is weaker than their (up to a factor $|V|$), it is more adapted to the structure of our graphs and has proven to be much faster in our experiments (see technical report [17]).

Some implementation details are crucial to the efficiency of the algorithm:

- **Exploiting connected components:** When there exists no arc between two subsets of V , it is possible to process them independently in order to solve the global min-cost flow problem.
- **Efficient max-flow algorithm:** We have implemented the “push-relabel” algorithm of [24] for solving our max-flow problems, using classical heuristics that significantly speed it up in practice (see [24, 26]). This algorithm leverages the concept of *pre-flow* that relaxes the definition of flow and allows vertices to have a positive excess. It can be initialized with any valid pre-flow, enabling warm-restarts when the max-flow is called several times as in our algorithm.
- **Improved projection step:** The first line of the function `computeFlow` can be replaced by $\gamma \leftarrow \arg \min_{\gamma} \sum_{j \in V_u} \frac{1}{2}(\mathbf{u}_j - \gamma_j)^2$ s.t. $\sum_{j \in V_u} \gamma_j \leq \lambda \sum_{g \in V_{gr}} \eta_g$ and $|\gamma_j| \leq \lambda \sum_{g \ni j} \eta_g$. The idea is that the structure of the graph will not allow $\bar{\xi}_j$ to be greater than $\lambda \sum_{g \ni j} \eta_g$ after the max-flow step. Adding these additional constraints leads to better performance when the graph is not well balanced. This modified projection step can still be computed in linear time [19].

3.3 Computation of the Dual Norm

The dual norm Ω^* of Ω , defined for any vector $\boldsymbol{\kappa}$ in \mathbb{R}^p by $\Omega^*(\boldsymbol{\kappa}) \triangleq \max_{\Omega(\mathbf{z}) \leq 1} \mathbf{z}^\top \boldsymbol{\kappa}$, is a key quantity to study sparsity-inducing regularizations [5, 15, 27]. We use it here to monitor the convergence of the proximal method through a duality gap, and define a proper optimality criterion for problem (1). We denote by f^* the Fenchel conjugate of f [28], defined by $f^*(\boldsymbol{\kappa}) \triangleq \sup_{\mathbf{z}} [\mathbf{z}^\top \boldsymbol{\kappa} - f(\mathbf{z})]$. The duality gap for problem (1) can be derived from standard Fenchel duality arguments [28] and it is equal to $f(\mathbf{w}) + \lambda\Omega(\mathbf{w}) + f^*(-\boldsymbol{\kappa})$ for $\mathbf{w}, \boldsymbol{\kappa}$ in \mathbb{R}^p with $\Omega^*(\boldsymbol{\kappa}) \leq \lambda$. Therefore, evaluating the duality gap requires to compute efficiently Ω^* in order to find a feasible dual variable $\boldsymbol{\kappa}$. This is equivalent to solving another network flow problem, based on the following variational formulation:

$$\Omega^*(\boldsymbol{\kappa}) = \min_{\boldsymbol{\xi} \in \mathbb{R}^{p \times |\mathcal{G}|}} \tau \quad \text{s.t.} \quad \sum_{g \in \mathcal{G}} \boldsymbol{\xi}^g = \boldsymbol{\kappa}, \text{ and } \forall g \in \mathcal{G}, \|\boldsymbol{\xi}^g\|_1 \leq \tau \eta_g \text{ with } \boldsymbol{\xi}_j^g = 0 \text{ if } j \notin g. \quad (5)$$

In the network problem associated with (5), the capacities on the arcs (s, g) , $g \in \mathcal{G}$, are set to $\tau \eta_g$, and the capacities on the arcs (j, t) , j in $[1; p]$, are fixed to $\boldsymbol{\kappa}_j$. Solving problem (5) amounts to finding the smallest value of τ , such that there exists a flow saturating the capacities $\boldsymbol{\kappa}_j$ on the arcs leading to the sink t (i.e., $\bar{\boldsymbol{\xi}} = \boldsymbol{\kappa}$). The algorithm below is proven to be correct in [17].

Algorithm 2 Computation of the dual norm.

- 1: **Inputs:** $\boldsymbol{\kappa} \in \mathbb{R}^p$, a set of groups \mathcal{G} , positive weights $(\eta_g)_{g \in \mathcal{G}}$.
- 2: Build the initial graph $G_0 = (V_0, E_0, s, t)$ as explained in Section 3.3.
- 3: $\tau \leftarrow \text{dualNorm}(V_0, E_0)$.
- 4: **Return:** τ (value of the dual norm).

Function $\text{dualNorm}(V = V_u \cup V_{gr}, E)$

- 1: $\tau \leftarrow (\sum_{j \in V_u} \boldsymbol{\kappa}_j) / (\sum_{g \in V_{gr}} \eta_g)$ and set the capacities of arcs (s, g) to $\tau \eta_g$ for all g in V_{gr} .
 - 2: Max-flow step: Update $(\bar{\boldsymbol{\xi}}_j)_{j \in V_u}$ by computing a max-flow on the graph (V, E, s, t) .
 - 3: **if** $\exists j \in V_u$ s.t. $\bar{\boldsymbol{\xi}}_j \neq \boldsymbol{\kappa}_j$ **then**
 - 4: Define (V^+, E^+) and (V^-, E^-) as in Algorithm 1, and set $\tau \leftarrow \text{dualNorm}(V^-, E^-)$.
 - 5: **end if**
 - 6: **Return:** τ .
-

4 Applications and Experiments

Our experiments use the algorithm of [4] based on our proximal operator, with weights η_g set to 1.

4.1 Speed Comparison

We compare our method (ProxFlow) and two generic optimization techniques, namely a subgradient descent (SG) and an interior point method,⁶ on a regularized linear regression problem. Both SG and ProxFlow are implemented in C++. Experiments are run on a single-core 2.8 GHz CPU. We consider a design matrix \mathbf{X} in $\mathbb{R}^{n \times p}$ built from overcomplete dictionaries of discrete cosine transforms (DCT), which are naturally organized on one- or two-dimensional grids and display local correlations. The following families of groups \mathcal{G} using this spatial information are thus considered: (1) every contiguous sequence of length 3 for the one-dimensional case, and (2) every 3×3 -square in the two-dimensional setting. We generate vectors \mathbf{y} in \mathbb{R}^n according to the linear model $\mathbf{y} = \mathbf{X}\mathbf{w}_0 + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, 0.01\|\mathbf{X}\mathbf{w}_0\|_2^2)$. The vector \mathbf{w}_0 has about 20% percent nonzero components, randomly selected, while respecting the structure of \mathcal{G} , and uniformly generated between $[-1, 1]$.

In our experiments, the regularization parameter λ is chosen to achieve the same sparsity as \mathbf{w}_0 . For SG, we take the step size to be equal to $a/(k+b)$, where k is the iteration number, and (a, b) are the best parameters selected in $\{10^{-3}, \dots, 10\} \times \{10^2, 10^3, 10^4\}$. For the interior point methods, since problem (1) can be cast either as a quadratic (QP) or as a conic program (CP), we show in Figure 2 the results for both formulations. Our approach compares favorably with the other methods, on three problems of different sizes, $(n, p) \in \{(100, 10^3), (1024, 10^4), (1024, 10^5)\}$, see Figure 2. In addition, note that QP, CP and SG do not obtain sparse solutions, whereas ProxFlow does. We have also run ProxFlow and SG on a larger dataset with $(n, p) = (100, 10^6)$: after 12 hours, ProxFlow and SG have reached a relative duality gap of 0.0006 and 0.02 respectively.⁷

⁶In our simulations, we use the commercial software Mosek, <http://www.mosek.com/>.

⁷Due to the computational burden, QP and CP could not be run on every problem.




Figure 2: Speed comparisons: distance to the optimal primal value versus CPU time (log-log scale).⁶




Figure 3: From left to right: original image \mathbf{y} ; estimated background $\mathbf{X}\mathbf{w}$; foreground (the sparsity pattern of \mathbf{e} used as mask on \mathbf{y}) estimated with ℓ_1 ; foreground estimated with $\ell_1 + \Omega$; another foreground obtained with Ω , on a different image, with the same values of λ_1, λ_2 as for the previous image. For the top row, the percentage of pixels matching the ground truth is 98.8% with Ω , 87.0% without. As for the bottom row, the result is 93.8% with Ω , 90.4% without (best seen in color).

4.2 Background Subtraction

Following [9, 10], we consider a background subtraction task. Given a sequence of frames from a fixed camera, we try to segment out foreground objects in a new image. If we denote by $\mathbf{y} \in \mathbb{R}^n$ a test image, we model \mathbf{y} as a sparse linear combination of p other images $\mathbf{X} \in \mathbb{R}^{n \times p}$, plus an error term \mathbf{e} in \mathbb{R}^n , i.e., $\mathbf{y} \approx \mathbf{X}\mathbf{w} + \mathbf{e}$ for some sparse vector \mathbf{w} in \mathbb{R}^p . This approach is reminiscent of [29] in the context of face recognition, where \mathbf{e} is further made sparse to deal with occlusions. The term $\mathbf{X}\mathbf{w}$ accounts for *background* parts present in both \mathbf{y} and \mathbf{X} , while \mathbf{e} contains specific, or *foreground*, objects in \mathbf{y} . The resulting optimization problem is $\min_{\mathbf{w}, \mathbf{e}} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{w} - \mathbf{e}\|_2^2 + \lambda_1 \|\mathbf{w}\|_1 + \lambda_2 \|\mathbf{e}\|_1$, with $\lambda_1, \lambda_2 \geq 0$. In this formulation, the ℓ_1 -norm penalty on \mathbf{e} does not take into account the fact that neighboring pixels in \mathbf{y} are likely to share the same label (background or foreground), which may lead to scattered pieces of foreground and background regions (Figure 3). We therefore put an additional structured regularization term Ω on \mathbf{e} , where the groups in \mathcal{G} are all the overlapping 3×3 -squares on the image. A dataset with hand-segmented evaluation images is used to illustrate the effect of Ω .⁸ For simplicity, we use a single regularization parameter, i.e., $\lambda_1 = \lambda_2$, chosen to maximize the number of pixels matching the ground truth. We consider $p = 200$ images with $n = 57600$ pixels (i.e., a resolution of 120×160 , times 3 for the RGB channels). As shown in Figure 3, adding Ω improves the background subtraction results for the two tested videos, by encoding, unlike the ℓ_1 -norm, both spatial and color consistency.

4.3 Multi-Task Learning of Hierarchical Structures

In [12], Jenatton et al. have recently proposed to use a hierarchical structured norm to learn dictionaries of natural image patches. Following this work, we seek to represent n signals $\{\mathbf{y}^1, \dots, \mathbf{y}^n\}$ of dimension m as sparse linear combinations of elements from a dictionary $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^p]$ in $\mathbb{R}^{m \times p}$. This can be expressed for all i in $[1; n]$ as $\mathbf{y}^i \approx \mathbf{X}\mathbf{w}^i$, for some sparse vector \mathbf{w}^i in \mathbb{R}^p . In [12], the dictionary elements are embedded in a *predefined* tree \mathcal{T} , via a particular instance of the structured norm Ω ; we refer to it as Ω_{tree} , and call \mathcal{G} the underlying set of groups. In this case, each signal \mathbf{y}^i admits a sparse decomposition in the form of a subtree of dictionary elements.

⁸<http://research.microsoft.com/en-us/um/people/jckrumm/wallflower/testimages.htm>

Inspired by ideas from multi-task learning [14], we propose to learn the tree structure \mathcal{T} by pruning irrelevant parts of a larger initial tree \mathcal{T}_0 . We achieve this by using an additional regularization term Ω_{joint} across the different decompositions, so that subtrees of \mathcal{T}_0 will *simultaneously* be removed for all signals \mathbf{y}^i . In other words, the approach of [12] is extended by the following formulation:

$$\min_{\mathbf{X}, \mathbf{W}} \frac{1}{n} \sum_{i=1}^n \left[\frac{1}{2} \|\mathbf{y}^i - \mathbf{X}\mathbf{w}^i\|_2^2 + \lambda_1 \Omega_{\text{tree}}(\mathbf{w}^i) \right] + \lambda_2 \Omega_{\text{joint}}(\mathbf{W}), \text{ s.t. } \|\mathbf{x}^j\|_2 \leq 1, \text{ for all } j \text{ in } [1; p], \quad (6)$$

where $\mathbf{W} \triangleq [\mathbf{w}^1, \dots, \mathbf{w}^n]$ is the matrix of decomposition coefficients in $\mathbb{R}^{p \times n}$. The new regularization term operates on the rows of \mathbf{W} and is defined as $\Omega_{\text{joint}}(\mathbf{W}) \triangleq \sum_{g \in \mathcal{G}} \max_{i \in [1; n]} |\mathbf{w}_g^i|$.⁹ The overall penalty on \mathbf{W} , which results from the combination of Ω_{tree} and Ω_{joint} , is itself an instance of Ω with general overlapping groups, as defined in Eq (2).

To address problem (6), we use the same optimization scheme as [12], i.e., alternating between \mathbf{X} and \mathbf{W} , fixing one variable while optimizing with respect to the other. The task we consider is the denoising of natural image patches, with the same dataset and protocol as [12]. We study whether learning the hierarchy of the dictionary elements improves the denoising performance, compared to standard sparse coding (i.e., when Ω_{tree} is the ℓ_1 -norm and $\lambda_2 = 0$) and the hierarchical dictionary learning of [12] based on predefined trees (i.e., $\lambda_2 = 0$). The dimensions of the training set — 50 000 patches of size 8×8 for dictionaries with up to $p = 400$ elements — impose to handle large graphs, with $|E| \approx |V| \approx 4.10^7$. Since problem (6) is too large to be solved many times to select the regularization parameters (λ_1, λ_2) rigorously, we use the following heuristics: we optimize mostly with the currently pruned tree held fixed (i.e., $\lambda_2 = 0$), and only prune the tree (i.e., $\lambda_2 > 0$) every few steps on a random subset of 10 000 patches. We consider the same hierarchies as in [12], involving between 30 and 400 dictionary elements. The regularization parameter λ_1 is selected on the validation set of 25 000 patches, for both sparse coding (Flat) and hierarchical dictionary learning (Tree). Starting from the tree giving the best performance (in this case the largest one, see Figure 4), we solve problem (6) following our heuristics, for increasing values of λ_2 . As shown in Figure 4, there is a regime where our approach performs significantly better than the two other compared methods. The standard deviation of the noise is 0.2 (the pixels have values in $[0, 1]$); no significant improvements were observed for lower levels of noise.

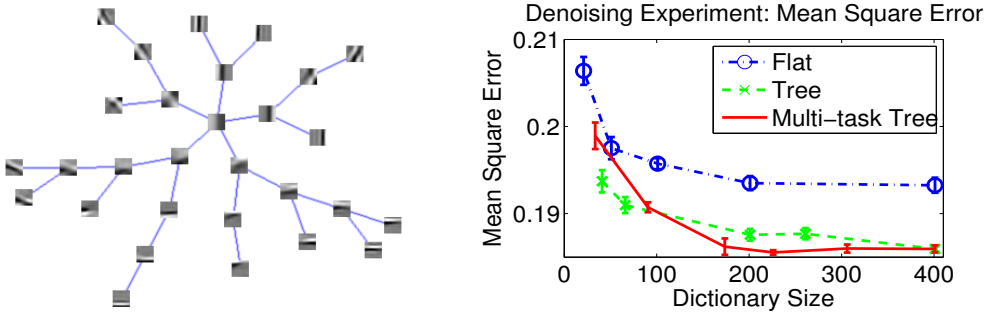


Figure 4: Left: Hierarchy obtained by pruning a larger tree of 76 elements. Right: Mean square error versus dictionary size. The error bars represent two standard deviations, based on three runs.

5 Conclusion

We have presented a new optimization framework for solving sparse structured problems involving sums of ℓ_∞ -norms of any (overlapping) groups of variables. Interestingly, this sheds new light on connections between sparse methods and the literature of network flow optimization. In particular, the proximal operator for the formulation we consider can be cast as a quadratic min-cost flow problem, for which we propose an efficient and simple algorithm. This allows the use of accelerated gradient methods. Several experiments demonstrate that our algorithm can be applied to a wide class of learning problems, which have not been addressed before within sparse methods.

Acknowledgments

This paper was partially supported by the European Research Council (SIERRA Project). The authors would like to thank Jean Ponce for interesting discussions and suggestions.

⁹The simplified case where Ω_{tree} and Ω_{joint} are the ℓ_1 - and mixed ℓ_1/ℓ_2 -norms [13] corresponds to [30].

References

- [1] P. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *Ann. Stat.*, 37(4):1705–1732, 2009.
- [2] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Ann. Stat.*, 32(2):407–499, 2004.
- [3] Y. Nesterov. Gradient methods for minimizing composite objective function. Technical report, Center for Operations Research and Econometrics (CORE), Catholic University of Louvain, 2007.
- [4] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imag. Sci.*, 2(1):183–202, 2009.
- [5] R. Jenatton, J-Y. Audibert, and F. Bach. Structured variable selection with sparsity-inducing norms. Technical report, 2009. Preprint arXiv:0904.3523v1.
- [6] L. Jacob, G. Obozinski, and J.-P. Vert. Group Lasso with overlap and graph Lasso. In *Proc. ICML*, 2009.
- [7] P. Zhao, G. Rocha, and B. Yu. The composite absolute penalties family for grouped and hierarchical variable selection. *Ann. Stat.*, 37(6A):3468–3497, 2009.
- [8] R. G. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE T. Inform. Theory*, 2010. to appear.
- [9] V. Cevher, M.F. Duarte, C. Hedge, and R.G. Baraniuk. Sparse signal recovery using markov random fields. In *Adv. NIPS*, 2008.
- [10] J. Huang, T. Zhang, and D. Metaxas. Learning with structured sparsity. In *Proc. ICML*, 2009.
- [11] S. Kim and E. P. Xing. Tree-guided group lasso for multi-task regression with structured sparsity. In *Proc. ICML*, 2010.
- [12] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for sparse hierarchical dictionary learning. In *Proc. ICML*, 2010.
- [13] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *J. Roy. Stat. Soc. B*, 68:49–67, 2006.
- [14] G. Obozinski, B. Taskar, and M. I. Jordan. Joint covariate selection and joint subspace selection for multiple classification problems. *Stat. Comput.*, 20(2):231–252, 2010.
- [15] F. Bach. Exploring large feature spaces with hierarchical multiple kernel learning. In *Adv. NIPS*, 2008.
- [16] P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*. Springer, 2010.
- [17] J. Mairal, R. Jenatton, G. Obozinski, and F. Bach. Network flow algorithms for structured sparsity. Technical report, 2010. Preprint arXiv:1008.5209v1.
- [18] D. S. Hochbaum and S. P. Hong. About strongly polynomial time algorithms for quadratic optimization over submodular constraints. *Math. Program.*, 69(1):269–309, 1995.
- [19] P. Brucker. An $O(n)$ algorithm for quadratic knapsack problems. *Oper. Res. Lett.*, 3:163–166, 1984.
- [20] G. Gallo, M. E. Grigoriadis, and R. E. Tarjan. A fast parametric maximum flow algorithm and applications. *SIAM J. Comput.*, 18:30–55, 1989.
- [21] M. Babenko and A.V. Goldberg. Experimental evaluation of a parametric flow algorithm. Technical report, Microsoft Research, 2006. MSR-TR-2006-77.
- [22] H. Groenevelt. Two algorithms for maximizing a separable concave function over a polymatroid feasible region. *Eur. J. Oper. Res.*, pages 227–236, 1991.
- [23] F. Bach. Structured sparsity-inducing norms through submodular functions. In *Adv. NIPS*, 2010.
- [24] A. V. Goldberg and R. E. Tarjan. A new approach to the maximum flow problem. In *Proc. of ACM Symposium on Theory of Computing*, 1986.
- [25] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian J. Math.*, 8(3), 1956.
- [26] B. V. Cherkassky and A. V. Goldberg. On implementing the pushrelabel method for the maximum flow problem. *Algorithmica*, 19(4):390–410, 1997.
- [27] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. In *Adv. NIPS*, 2009.
- [28] J. M. Borwein and A. S. Lewis. *Convex analysis and nonlinear optimization*. Springer, 2006.
- [29] J. Wright, A. Y. Yang, A. Ganesh, S. S. Sastry, and Y. Ma. Robust face recognition via sparse representation. *IEEE T. Pattern. Anal.*, pages 210–227, 2008.
- [30] P. Sprechmann, I. Ramirez, G. Sapiro, and Y. C. Eldar. Collaborative hierarchical sparse modeling. Technical report, 2010. Preprint arXiv:1003.0400v1.