
Combinatorial Preconditioners for Proximal Algorithms on Graphs

Thomas Möllenhoff

Zhenzhang Ye

Tao Wu

Daniel Cremers

Department of Computer Science, Technical University of Munich, Germany
{thomas.moellenhoff, zhenzhang.ye, tao.wu, cremers}@tum.de

Abstract

We present a novel preconditioning technique for proximal optimization methods that relies on graph algorithms to construct effective preconditioners. Such combinatorial preconditioners arise from partitioning the graph into forests. We prove that certain decompositions lead to a theoretically optimal condition number. We also show how ideal decompositions can be realized using matroid partitioning and propose efficient greedy variants thereof for large-scale problems. Coupled with specialized solvers for the resulting scaled proximal subproblems, the preconditioned algorithm achieves competitive performance in machine learning and vision applications.

1 Introduction

Many applications in statistics [48], learning [27], and imaging [11] rely on efficiently solving convex-concave saddle-point problems:

$$\max_{p \in \mathbb{R}^{\mathcal{E}}} \min_{u \in \mathbb{R}^{\mathcal{V}}} G(u) - F^*(p) + \langle Ku, p \rangle. \quad (1)$$

Here the model is formulated on an undirected weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \omega)$, whose edges are weighted by a given function $\omega : \mathcal{E} \rightarrow \mathbb{R}_+$. The extended real-valued functions $F : \mathbb{R}^{\mathcal{E}} \rightarrow \mathbb{R} \cup \{+\infty\}$ and $G : \mathbb{R}^{\mathcal{V}} \rightarrow \mathbb{R} \cup \{+\infty\}$ are assumed to be proper, lower semi-continuous and convex. The notation F^* refers to the convex conjugate of F . The (linear) vertex-to-edge map $K : \mathbb{R}^{\mathcal{V}} \rightarrow \mathbb{R}^{\mathcal{E}}$ is defined by

$$K = \text{diag}(\omega)\nabla,$$

where ∇ is the (transposed) incidence matrix of \mathcal{G} , i.e.

$$(\nabla u)_e = u_i - u_j, \quad \forall e = (i, j) \in \mathcal{E},$$

with arbitrarily fixed orientation. For F being the ℓ^1 -norm, i.e. $F(\cdot) = \|\cdot\|_1$, this choice yields the total-variation semi-norm of functions on a weighted graph. In addition to its ubiquitous applications in image processing and computer vision, this semi-norm has recently gained considerable attention in unsupervised learning [27, 28, 8], semi-supervised learning [22], collaborative filtering [6], clustering [22] and statistical inference [52].

Among other proximal algorithms (see [41, 13] and the references therein for an overview), the primal-dual hybrid gradient (PDHG) algorithm [2, 55, 43, 19, 11] is a popular solver for the problem in (1). A general formulation of PDHG iterations [13] appears as follows:

$$\begin{aligned} u^{k+1} &= \arg \min_{u \in \mathbb{R}^{\mathcal{V}}} G(u) + \langle p^k, Ku \rangle + \frac{s}{2} \|u - u^k\|_S^2, \quad (2) \\ p^{k+1} &= \arg \min_{p \in \mathbb{R}^{\mathcal{E}}} F^*(p) - \langle K(2u^{k+1} - u^k), p \rangle \\ &\quad + \frac{t}{2} \|p - p^k\|_T^2. \quad (3) \end{aligned}$$

Here $S \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ and $T \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{E}|}$ are symmetric positive definite matrices, such that $\|\cdot\|_S$ is a scaled norm defined by $\|u\|_S^2 = \langle u, u \rangle_S = u^\top S u$ and analogously for $\|\cdot\|_T$. For given S and T , the convergence of PDHG is guaranteed if the (inverse) step sizes s, t satisfy $st > \|T^{-1/2} K S^{-1/2}\|^2$, cf. [42, Lemma 1]. Interestingly, as pointed out by [19, 11, 13], the formulation in (2)–(3) provides a flexible framework for deriving various types of proximal splitting algorithms, e.g. the proximal gradient method, Douglas-Rachford splitting, and (linearized) ADMM. Proximal algorithms in form of (2)–(3) are particularly efficient when F, G admit *separable* structures and S, T are *diagonal*. In this scenario, solutions of the subproblems in (2) and (3) refer to pointwise proximal evaluations. Nonetheless, many research efforts have been devoted to further accelerating the convergence speed.

To this end, two categories of acceleration strategies are envisaged: multi-step acceleration and preconditioning. Classical (optimal) multi-step gradient descent methods are attributed to [44, 38]. In the context of proximal algorithms, the FISTA algorithm [4] was proposed as an accelerated proximal gradient method, and in [12] a multi-step PDHG was devised.

On the other hand, a preconditioning technique aims to accelerate convergence, typically through reducing the number of outer iterations, by choosing proper scaling matrices S and T (also called the preconditioners in this context). In contrast to their counterparts for solving linear systems, preconditioning techniques for proximal algorithms are much less developed. Diagonal preconditioners for PDHG were explored in [42]; Preconditioners for other types of proximal algorithms [5, 23, 24, 35, 54, 7, 25, 20] also appeared recently in the literature.

A consensus among all existing preconditioning approaches appears that, while favorably reducing the number of outer iterations, preconditioning could explode the overall computational load by expensive inner proximal evaluations. For example, for proximal gradient method in minimizing a sum of smooth and nonsmooth functions, using the (approximate) Hessian of the smooth function as preconditioner may attain superlinear convergence in the outer iteration but also lead to very expensive proximal (or backward) steps; see [35]. This is the reason why diagonal preconditioners remain a popular choice in many recent works, and non-diagonal preconditioners designed in [5, 54, 21] do not deviate far from the diagonal ones.

In this work, we propose combinatorial preconditioners for proximal algorithms based on a partitioning of the original graph into forests. This leads to a class of block diagonal preconditioners, and the resulting PDHG updates refer to solving parallel subproblems on forests. We show how to construct such preconditioners guided by theoretical estimates of the condition number. Coupled with fast direct solvers for proximal evaluation on forests, we achieve significant performance boost for the PDHG algorithm across a series of numerical tests.

2 Preconditioner and Condition Number

The choice of S and T can significantly influence the convergence speed of the (generalized) PDHG scheme, (2)–(3), in practice. In [42], Pock and Chambolle showed that utilization of diagonal preconditioners S and T yields a visible performance boost in comparison with PDHG without preconditioning (i.e. $S = I$, $T = I$). In a slightly different context, Boyd et al. [16, 20, 23, 24, 25] also considered diagonal preconditioning strategies for other closely related proximal algorithms. In particular, they suggested based on extensive numerical experiments that an ideal choice of S and T ought to minimize the (*finite*) *condition number* $\kappa(T^{-1/2}KS^{-1/2})$ defined by

$$\kappa(\cdot) := \frac{\sigma_{\max}(\cdot)}{\sigma_{\min>0}(\cdot)}, \quad (4)$$

i.e. the ratio between largest and smallest non-zero singular value. This rule of thumb was computationally pursued by so called matrix equilibration [16, 25, 1].

A more quantified connection between the convergence rate of (2)–(3) and $\kappa(T^{-1/2}KS^{-1/2})$ can be drawn in a more specific setting, e.g. $G(u) = \frac{1}{2}\|u - f\|^2$ for some given $f \in \mathbb{R}^V$. By choosing $s = 1$, $\tilde{S} = I$ in (2)–(3), one comes up with the following proximal gradient iteration:

$$p^{k+1} = \arg \min_{p \in \mathbb{R}^E} F^*(p) + \frac{t}{2} \|p - (p^k - (tT)^{-1}K(K^\top p^k - f))\|_T^2. \quad (5)$$

When $F^* = \delta_Q$ is the indicator function of a polyhedral set $Q := \{p \in \mathbb{R}^E : Ap \leq b\}$ it can be shown that iteration (5) converges linearly [37]. More precisely, with step size choice $t = \sigma_{\max}(T^{-1/2}K)$ the linear convergence rate $r \in [0, 1)$ can be derived (cf. [37, Theorem 11]) as:

$$r = \frac{\kappa^2 - 1}{\kappa^2 + 1}, \quad \kappa = \theta(T^{-1/2}A, T^{-1/2}K)\sigma_{\max}(T^{-1/2}K),$$

where $\theta(\cdot, \cdot)$ denotes Hoffman’s bound [29, 50]. In the extreme case where $F^* \approx 0$ (e.g., for strong regularizations), Hoffman’s bound reduces to $\sigma_{\min>0}(T^{-1/2}K)^{-1}$ and κ matches the condition number $\kappa(T^{-1/2}K)$.

In a general setting, a strong correlation between the convergence speed and the condition number is supported by the numerical evidences in Section 4.1 and Table 1.

3 Combinatorial Preconditioners

Following the discussion in the previous section, reducing $\kappa(T^{-1/2}KS^{-1/2})$ provides a reasonable guideline for the choice of preconditioners. As K is the weighted incidence matrix of \mathcal{G} , this boils down to constructing “good” approximations of the graph. Meanwhile, as discussed in the introduction, the non-diagonal S or T could possibly explode the computational cost of the proximal evaluation in (2)–(3). Hence, an ideal choice of preconditioners would strike a balance between reduced (outer) iteration number and costlier proximal evaluations per iteration.

Bearing this in mind, in this section we propose our combinatorial preconditioners based on graph partitioning. This leads to a family of block diagonal preconditioners for T . For simplicity we fix $S = I$ in our development, and remark that T can be adapted to work properly with any diagonal preconditioner S (e.g. the one used in [42]). In terms of the update scheme (3), such preconditioning yields proximal evaluation on respective partitioned subgraphs, which can be efficiently carried out by using the state-of-the-art direct solver on trees [34].

As a remark, there is a connection between our combinatorial preconditioners and the subgraph preconditioners for solving linear systems in graph Laplacians; see, e.g., [46] and the references therein. Pioneered by Vaidya and his coworkers in the early 1990s [49, 32], a series of works

are done in finding a subgraph preconditioner. An ideal subgraph preconditioner uses a graph Laplacian on some (low-stretch) spanning tree which best preserves connectivity between vertices in the original graph. In some sense, our proposed combinatorial preconditioners are dual analogues of the subgraph preconditioners.

In Section 3.1, we show how to construct T via graph partitioning and under which sufficient conditions the condition number $\kappa(T^{-1/2}K)$ can be (optimally) reduced. In Section 3.2, such sufficient conditions are algorithmically realized by: (1) chains on regular grid; (2) nested forests on general graph. In Section 3.3, we detail the message passing based implementation of efficient proximal evaluation on forests.

3.1 Preconditioning via Graph Partitioning

Let the edge set \mathcal{E} be partitioned into L mutually disjoint subsets, i.e. $\mathcal{E} = \bigsqcup_{l=1}^L \mathcal{E}_l$, such that each subgraph $\mathcal{G}_l = (\mathcal{V}, \mathcal{E}_l, \omega|_{\mathcal{E}_l})$ is a *forest*, i.e. \mathcal{G}_l has no cycle. Correspondingly, we define P_l as the canonical projection from $\mathbb{R}^{\mathcal{E}}$ to $\mathbb{R}^{\mathcal{E}_l}$, i.e. $P_l p = p|_{\mathcal{E}_l}$ for each $p \in \mathbb{R}^{\mathcal{E}}$. Thus, the matrix K can be decomposed into submatrices $\{K_l\}_{l=1}^L$ where each $K_l = P_l K \in \mathbb{R}^{|\mathcal{E}_l| \times |\mathcal{V}|}$. Analogously let $\nabla_l = P_l \nabla$ and $\omega_l = \omega|_{\mathcal{E}_l}$. Note that each ∇_l^\top has full column rank.

We then define our preconditioners as follows

$$T_l := K_l K_l^\top \quad \forall l \in \{1, \dots, L\},$$

$$T := \sum_{l=1}^L P_l^\top T_l P_l.$$

It follows from (4) that

$$\kappa(T^{-1/2}K) = \sqrt{\frac{\lambda_{\max}(\Pi)}{\lambda_{\min>0}(\Pi)}}, \quad (6)$$

$$\Pi := K^\top T^{-1} K = \sum_{l=1}^L \Pi_l, \quad \Pi_l := \nabla_l^\top (\nabla_l \nabla_l^\top)^{-1} \nabla_l. \quad (7)$$

Indeed, each Π_l is the orthogonal projection onto the subspace $\text{ran } \nabla_l^\top$, and hence $\text{ran } \Pi_l = \text{ran } \nabla_l^\top$ and $\ker \Pi_l = \ker \nabla_l$. It follows immediately from (7) that

$$\lambda_{\max}(\Pi) \leq L, \quad (8)$$

and therefore it suffices for the sake of convergence to choose step sizes s, t such that $st > L$.

In the remainder of Section 3.1, we show analytically how certain graph partitions can attain optimal condition number $\kappa(T^{-1/2}K)$ in a two-partition scenario, i.e. $L = 2$. The two-partition scenario is motivated from many applications where \mathcal{G} is a 2D regular grid (whose maximum degree equals 4). In this case, it is guaranteed by Nash-Williams' theorem [36] that \mathcal{G} can be covered by two disjoint forests. Our theory suggests that either a *chain* decomposition or a

nested-forest decomposition makes a good preconditioner on the 2D regular grid. Furthermore, the nested-forest preconditioners further extend to any general weighted graph, with aid of either the matroid partition algorithm [18] or some greedy algorithm; see Section 3.2.2.

The following theorem derives a lower bound for $\kappa(T^{-1/2}K)$ for a wide range of two-partition cases. We denote by $\text{rank } \nabla^\top$ the column rank of ∇^\top . In terms of graph theory, $\text{rank } \nabla^\top$ identifies the maximum number of the edges in \mathcal{G} that are cycle-free. T^{-1} is understood as the Moore-Penrose pseudoinverse when T is singular. The two assumptions rule out pathological cases where \mathcal{G} is either too sparse (e.g. cycle-free) or too dense (e.g. fully connected). In the proof, we shall use Weyl's inequality [51]:

$$\lambda_{i+j-1}(\Pi_1 + \Pi_2) \leq \lambda_i(\Pi_1) + \lambda_j(\Pi_2), \quad (9)$$

where $i, j \geq 1$, $i + j - 1 \leq |\mathcal{V}|$, and $\lambda_i(\cdot)$ denotes the i -th largest eigenvalue of a real symmetric matrix.

Theorem 1. *Let \mathcal{G} be partitioned into two nonempty subgraphs \mathcal{G}_1 and \mathcal{G}_2 (not necessarily forests) such that*

1. $\text{ran } \nabla_1^\top \cap \text{ran } \nabla_2^\top \supseteq \{0\}$.
2. $\text{rank } \nabla^\top > \min(\text{rank } \nabla_1^\top, \text{rank } \nabla_2^\top)$.

Then we have $\kappa(T^{-1/2}K) \geq \sqrt{2}$.

Proof. By Weyl's inequality (9), we have $\lambda_1(\Pi) \leq \lambda_1(\Pi_1) + \lambda_1(\Pi_2) = 2$. Moreover, condition 1 ensures that there exists some nonzero $v \in \mathbb{R}^{\mathcal{V}}$ with $\Pi v = \Pi_1 v + \Pi_2 v = 2v$. Hence, we must have $\lambda_1(\Pi) = 2$.

Let $r = \text{rank } \nabla^\top$. Without loss of generality, we assume $\text{rank } \nabla_1^\top \geq \text{rank } \nabla_2^\top$. Hence $r > \text{rank } \nabla_2^\top$ due to condition 2. Again using (9), we have $\lambda_{\min>0}(\Pi) = \lambda_r(\Pi) \leq \lambda_1(\Pi_1) + \lambda_r(\Pi_2) = 1 + 0 = 1$. Altogether, the conclusion follows in view of (6). \square

We proceed to propose sufficient conditions for graph partitioning (in particular when $L = 2$) which guarantee optimal condition number $\kappa(T^{-1/2}K) = \sqrt{2}$.

Theorem 2. *Let \mathcal{G} be partitioned into two forests \mathcal{G}_1 and \mathcal{G}_2 . If the following conditions are satisfied*

1. $\Pi_1 \Pi_2 = \Pi_2 \Pi_1$,
2. $\ker(\Pi_1 \Pi_2) \setminus \ker \Pi \neq \emptyset$,
3. $\text{ran } \Pi_1 \cap \text{ran } \Pi_2 \supseteq \{0\}$,

then $\kappa(T^{-1/2}K) = \sqrt{2}$.

Proof. The proof of $\lambda_{\max}(\Pi) = 2$ is identical to that for Theorem 1.

On the other hand, we have

$$\begin{aligned}\lambda_{\min>0}(\Pi) &= \min\{\langle u, \Pi u \rangle / \|u\|^2 : u \in \text{ran } \Pi, u \neq 0\}, \\ &= \min\{\langle v, \Pi^3 v \rangle / \|\Pi v\|^2 : \Pi v \neq 0\},\end{aligned}$$

where by commutativity (i.e. condition 1)

$$\begin{aligned}\frac{\langle v, \Pi^3 v \rangle}{\|\Pi v\|^2} &= \frac{\langle v, \Pi_1 v \rangle + \langle v, \Pi_2 v \rangle + 6 \langle v, \Pi_1 \Pi_2 v \rangle}{\langle v, \Pi_1 v \rangle + \langle v, \Pi_2 v \rangle + 2 \langle v, \Pi_1 \Pi_2 v \rangle} \\ &= 1 + \frac{4 \langle v, \Pi_1 \Pi_2 v \rangle}{\langle v, \Pi_1 v \rangle + \langle v, \Pi_2 v \rangle + 2 \langle v, \Pi_1 \Pi_2 v \rangle}.\end{aligned}$$

Note that $\Pi_1 \Pi_2$ is an orthogonal projection onto $\text{ran } \Pi_1 \cap \text{ran } \Pi_2$, and hence $\langle v, \Pi_1 \Pi_2 v \rangle \geq 0$ for all v . Meanwhile, condition 2 asserts that $\Pi v \neq 0$ and $\langle v, \Pi_1 \Pi_2 v \rangle = 0$ for some v . Altogether, $\lambda_{\min>0}(\Pi) = 1$ and the conclusion follows. \square

Theorem 3. *Let \mathcal{G} be partitioned into L (nonempty) nested forests, namely $\{\mathcal{G}_l\}_{l=1}^L$, in the sense that*

$$\text{ran } \Pi_1 = \dots = \text{ran } \Pi_{\widehat{l}} \supseteq \text{ran } \Pi_{\widehat{l}+1} \supseteq \dots \supseteq \text{ran } \Pi_L.$$

Then we have $\kappa(T^{-1/2}K) = \sqrt{L/\widehat{l}}$.

Proof. Recall (8) that $\lambda_{\max}(\Pi) \leq L$. Indeed, we have $\lambda_{\max}(\Pi) = L$ since $\Pi v = Lv$ for any nonzero $v \in \text{ran } \Pi_L$. On the other hand, note that $\text{ran } \Pi = \text{ran } \Pi_1 = \dots = \text{ran } \Pi_{\widehat{l}}$. Therefore, for all $v \in \text{ran } \Pi$, we have $\langle v, \Pi v \rangle \geq \sum_{l=1}^{\widehat{l}} \langle v, \Pi_l v \rangle = \widehat{l} \|v\|^2$ and the equality holds when $v \in \text{ran } \Pi \setminus \text{ran } \Pi_{\widehat{l}+1}$. This gives $\lambda_{\min>0}(\Pi) = \widehat{l}$, which concludes the proof. \square

3.2 Two Classes of Forest Preconditioners

Given the abstract conditions from Theorem 2 and 3 under which graph partitioning achieves optimal condition number, the question remains whether such partitions exist and how they can be constructed in practice. In the following section we propose optimal partitioning approaches, depending on the topology of the underlying graph. We first focus on regular grid graphs, which are ubiquitous in signal processing and computer vision applications.

3.2.1 Chains on Regular Grid

Formally, d -dimensional grid graphs are given as the d -fold Cartesian product $\mathcal{G} = \mathcal{P}_1 \square \dots \square \mathcal{P}_d$ between path graphs $\{\mathcal{P}_l\}_{l=1}^d$, $\mathcal{P}_l = (\{1, \dots, n_l\}, \{(i, i+1)\}_{i=1}^{n_l-1})$. For ease of presentation we will focus mostly on two dimensional grids. Recall that the Cartesian product $(\mathcal{V}_1, \mathcal{E}_1) \square (\mathcal{V}_2, \mathcal{E}_2)$ yields a graph with vertex set $\mathcal{V}_1 \times \mathcal{V}_2$. In that graph, two vertices (v_1, v_2) and (v'_1, v'_2) are adjacent if and only if $v_1 = v'_1$ and v_2 is adjacent to v'_2 or $v_2 = v'_2$ and v_1 is adjacent to v'_1 . Whether a given graph

is a grid graph can be tested in linear time [30]. The procedure described in [30] yields a natural decomposition of the d -dimensional grid into d linear forests, i.e., a forest in which every tree is a path graph. Moreover, in the aforementioned decomposition, each forest consists of paths of equal length. To be precise, in the case $d = 2$ the decomposition is given by

$$\mathcal{G}_1 = \mathcal{I}_{n_2} \square \mathcal{P}_1, \quad \mathcal{G}_2 = \mathcal{P}_2 \square \mathcal{I}_{n_1}, \quad (10)$$

where $\mathcal{I}_n = (\{1, \dots, n\}, \emptyset)$. We refer to the connected components of \mathcal{G}_1 as horizontal chains and \mathcal{G}_2 as vertical chains. Such a natural splitting into chains enjoys a theoretically optimal condition number among the large class of decompositions characterized by Theorem 1.

Theorem 4. *Let $\mathcal{G} = \mathcal{P}_1 \square \mathcal{P}_2$ be a regular grid of dimension $n_1 \times n_2$ ($n_1, n_2 \geq 3$). The condition number in the unpreconditioned case is bounded by*

$$\kappa(K) \geq \sqrt{2} \max(n_1, n_2) / \sqrt{\pi^2/2}. \quad (11)$$

The natural partitioning of \mathcal{G} into two linear forests (10) achieves condition number

$$\kappa(T^{-1/2}K) = \sqrt{2}. \quad (12)$$

Proof. The bound (11) follows from elementary results in spectral graph theory. The spectrum $\{\lambda_i\}$ of a path graph of dimension n is $\lambda_i = 2 - 2 \cos(\pi i/n)$, $i \in \{1, \dots, n-1\}$, see [9, Section 1.4.4]. The spectrum of the Cartesian product graph is given by the summing all pairs of eigenvalues of the individual graphs, cf. [9, Section 1.4.6]. Using the inequality $1 - \cos(x) \leq x^2/2$ we have

$$\kappa(K)^2 = \frac{2 - \sum_{l=1}^2 \cos\left(\frac{\pi(n_l-1)}{n_l}\right)}{1 - \cos\left(\frac{\pi}{\max(n_1, n_2)}\right)} \geq \frac{2 \max(n_1, n_2)^2}{\pi^2/2}.$$

To show (12) it suffices to verify the three conditions in Theorem 2. Denoting by $\mathbf{1}_n$ the ones-vector and by \otimes the Kronecker product, the projections are explicitly given by

$$\begin{aligned}\Pi_1 &= I_{n_2} \otimes \left(I_{n_1} - \frac{1}{n_1} (\mathbf{1}_{n_1} \mathbf{1}_{n_1}^\top) \right), \\ \Pi_2 &= \left(I_{n_2} - \frac{1}{n_2} (\mathbf{1}_{n_2} \mathbf{1}_{n_2}^\top) \right) \otimes I_{n_1}.\end{aligned}$$

Since $(A \otimes B)(C \otimes D) = (AC \otimes BD)$, the condition $\Pi_1 \Pi_2 = \Pi_2 \Pi_1$ from Theorem 2 follows directly.

Now note that each projection Π_i subtracts the mean on the chains it contains. Conditions 2 and 3 can be verified by counting dimensions. For condition 2, simply pick nonzero $u \in \mathbb{R}^{\mathcal{V}}$ to be constant along chains in \mathcal{P}_1 but non-constant along chains in \mathcal{P}_2 . For condition 3, pick nonzero $u \in \mathbb{R}^{\mathcal{V}}$ that is zero-mean on the chains in both \mathcal{P}_1 and \mathcal{P}_2 . \square

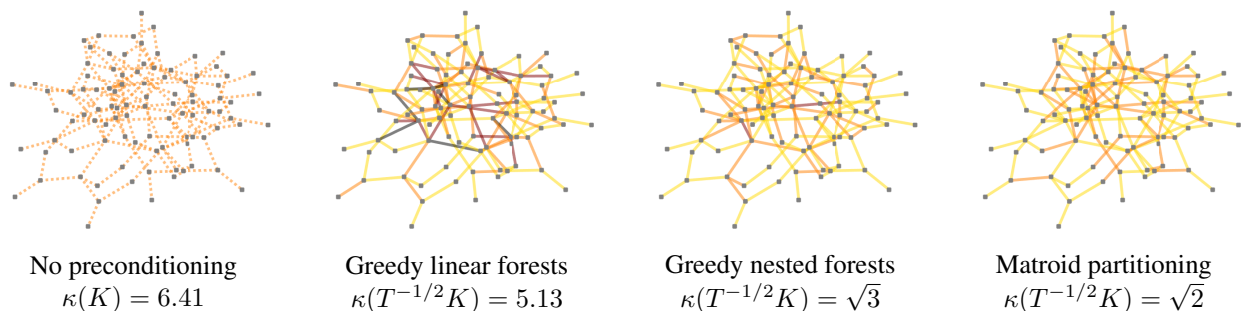


Figure 1: Overview of the proposed graph decompositions on an example graph ($|\mathcal{V}| = 96$, $|\mathcal{E}| = 150$). The decomposition into linear forests does not satisfy the assumptions of Theorem 3 and does not lead to a significant reduction in condition number. While the greedy nested forest decomposition satisfies the assumptions, it finds a partition into three spanning forests leading to a suboptimal condition number. The matroid approach guarantees the best possible condition number.

We remark that the above results can be generalized to the d -dimensional setting, yielding $\kappa(K) \geq \sqrt{d}(\max_i n_i)/\sqrt{\pi^2/2}$ and $\kappa(T^{-1/2}K) = \sqrt{d}$. Remarkably the chains preconditioning mentioned above makes the condition number independent of the grid size, while for the unpreconditioned case, it exhibits linear growth with respect to the largest grid dimension.

Furthermore, the splitting into chains leads to a particularly efficient evaluation of the dual subproblem in PDHG as we will see later on. Together with the theoretically optimal condition number, this makes the chains the number one choice of preconditioner for the regular d -dimensional grid.

3.2.2 Nested Forests on General Graphs

Matroid partitioning. The situation on general graphs is more involved than on grid graphs. Inspired by Theorem 3, we seek to partition \mathcal{G} into nested forests such that the number of forests is *minimal*, i.e. equal to the *arboricity* of the graph [36]. If \mathcal{G} is connected, the arboricity of \mathcal{G} can be calculated as: $\max \left\{ \left\lceil \frac{|\mathcal{E}_{\mathcal{G}'}}{|\mathcal{V}_{\mathcal{G}'}|-1} \right\rceil : \mathcal{V}_{\mathcal{G}'} \subset \mathcal{V}_{\mathcal{G}}, \mathcal{E}_{\mathcal{G}'} \subset \mathcal{E}_{\mathcal{G}} \right\}$.

It turns out that the classical matroid partitioning algorithm by Edmonds [18] meets our requirements. In short, matroid partitioning progressively inserts an idle edge into the forest partitions. To preserve all partitions cycle-free, it relies on a primitive operation which detects a simple cycle (also called *circuit*) whenever this occurs due to the insertion of a new edge into a forest. By the nature of the algorithm, the resulting partitions are guaranteed to be (a minimal number of) nested forests, and hence realize the condition posed in Theorem 3.

In spite of such favorable properties of matroid partitioning, its complexity grows like $\mathcal{O}(|\mathcal{E}|^3 + |\mathcal{E}|^2L)$, cf. [33], making its application prohibitive for large graphs.

Greedy nested forests. As a remedy, we propose the following “greedy nested forests” heuristic: given the input graph \mathcal{G} we successively subtract a spanning forest until

$\frac{ \mathcal{E} }{ \mathcal{V} }$	No Precond.		Linear Forest		Nested Forest			
	κ^2	it	κ^2	it	Greedy		Matroid	
	κ^2	it	κ^2	it	κ^2	it	κ^2	it
0.68	53.8	1624	24.9	789	1	36	1	36
1.11	257.0	6245	148.7	3782	2	80	2	80
2.36	82.2	2061	14.1	577	4	143	3	103
2.98	27.7	1010	12.4	484	5	194	4	137
4.25	24.9	800	10.1	419	7	254	5	166
6.48	7.7	367	5.6	176	10	362	1.75	59

Table 1: Comparison of condition number and PDHG iterations for various forest strategies on small random graphs ($|\mathcal{V}| = 512$) with varying edge to vertex ratio $\frac{m}{n}$.

no edges remain. The individual subtracted forests form the graph partitioning $\{\mathcal{G}_l\}_{l=1}^L$. While this greedy approach does not guarantee a minimal number of forests L , the partition still satisfies the assumption of Theorem 3. Indeed, each edge in the forest \mathcal{G}_l can be represented by a path in \mathcal{G}_{l-1} since adding that edge from \mathcal{G}_l to \mathcal{G}_{l-1} would form a cycle due to the spanning forest property.

Greedy linear forests. Since the dual update can be computed very efficiently for linear forests we modify the above procedure to yield a linear forest decomposition. Before subtracting the spanning forest from the residual graph, we remove all edges which contain a vertex with degree larger than 2. In addition, we check whether it is possible to add any edges from the residual graph to the current linear forest without turning it into a general forest.

In Table 1 we show the condition number for the different partitioning strategies on small random Erdős-Renyi graphs of varying average degree. While the matroid partitioning strategy finds the lowest condition number, both greedy heuristics also lead to a reduction in condition number for most cases. The greedy nested forest heuristic works best for graphs with low edge-to-vertex ratio, while the linear forest heuristic is preferable for dense graphs.

Algorithm 1 Total variation on a forest [34, Algorithm 2].

Input: Weighted forest $\mathcal{G}_l = (\mathcal{V}, \mathcal{E}_l, \omega_l)$, $f_l \in \mathbb{R}^{\mathcal{V}}$.

for each tree $\mathcal{T} = (\mathcal{V}', \mathcal{E}', \omega')$ in \mathcal{G}_l **do**
 // Message passing from leaves to the root.
for each $(i, j) \in \mathcal{E}'$ from leaves to root $r \in \mathcal{V}'$ **do**
 $\widehat{m}_i(u) = u - f_{l,i} + \sum_{(j,i) \in \mathcal{E}'}$ $\widehat{m}_{j,i}(u)$.
 $m_{i,j}(u) = \text{clip}_{[-\omega'_{i,j}, \omega'_{i,j}]}(\widehat{m}_i(u))$.
 find $\lambda_{i,j}^-, \lambda_{i,j}^+$ which satisfy both
 $\widehat{m}_i(\lambda_{i,j}^-) = -\omega'_{i,j}$, $\widehat{m}_i(\lambda_{i,j}^+) = \omega'_{i,j}$.
end for
 // Compute solution on tree.
 Solve $\widehat{m}_r(v_{l,r}) = 0$ for $v_{l,r}$.
for each $(i, j) \in \mathcal{E}'$ from root toward leaves **do**
 $v_{l,i} = \text{clip}_{[\lambda_{i,j}^-, \lambda_{i,j}^+]}(v_{l,j})$.
end for
end for
Output: $v_l \in \mathbb{R}^{\mathcal{V}}$.

3.3 Proximal Evaluation on Forests

In this section we will discuss how the dual update (3) is computed for our combinatorial preconditioner T . Assuming that F^* is separable across the subgraphs $\{\mathcal{G}_l\}_{l=1}^L$,

$$F^*(p) = \sum_{l=1}^L F_l^*(p|_{\mathcal{E}_l}),$$

the dual update (3) decomposes into parallel problems:

$$p^{k+1}|_{\mathcal{E}_l} = \arg \min_{p \in \mathbb{R}^{\mathcal{E}_l}} \frac{t}{2} \|p - p^k|_{\mathcal{E}_l}\|_{T_l}^2 + F_l^*(p) - \langle K \bar{u}^{k+1}|_{\mathcal{E}_l}, p \rangle, \quad \forall 1 \leq l \leq L. \quad (13)$$

Let us now further assume that the individual \mathcal{G}_l are forests and $F(\cdot) = \|\cdot\|_1$.

Expanding the norm $\|\cdot\|_{T_l}^2$ and completing the square in (13) leads to the problem

$$p^{k+1}|_{\mathcal{E}_l} = \arg \min_{\|p\|_{\infty} \leq 1} \frac{1}{2} \|K_l^\top p + f_l\|^2, \quad (14)$$

with $f_l = -K_l^\top(p^k|_{\mathcal{E}_l}) - (2u^{k+1} - u^k)/t$. The dual problem to (14) is given by

$$v_l = \arg \min_{u \in \mathbb{R}^{\mathcal{V}}} \frac{1}{2} \|u - f_l\|^2 + \|K_l u\|_1. \quad (15)$$

This further parallelizes into weighted total variation problems on the individual trees in the forest \mathcal{G}_l . These problems can be handled due to recent advances in direct total-variation solvers; see [15, 17, 31, 14, 3, 34]. The original taut-string algorithm [15, 17] solves the 1D total-variation problem in $\mathcal{O}(n)$ iterations on a chain. Condat [14] proposed an algorithm which has worst-case $\mathcal{O}(n^2)$ complexity but it achieves good performance in practice. Barbero

Algorithm 2 PDHG with combinatorial preconditioning for total variation minimization on weighted graphs.

Input: $u^0 \in \mathbb{R}^{\mathcal{V}}$, $p^0 \in \mathbb{R}^{\mathcal{E}}$, $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \omega)$.

Compute decomposition of \mathcal{G} into forests $\{\mathcal{G}_l\}_{l=1}^L$.

Pick $s, t > 0$ satisfying $st > L$.

for $k \geq 0$ **while** not converged **do**
 // primal update
 $u^{k+1} = \arg \min_u \frac{s}{2} \|u - u^k\|^2 + \langle K^\top p^k, u \rangle + G(u)$.
 $\bar{u}^{k+1} = 2u^{k+1} - u^k$.
 // dual update
for each forest $l = 1 \dots L$ **do**
 $f_l = -K_l^\top(p^k|_{\mathcal{E}_l}) - \bar{u}^{k+1}/t$.
 Obtain v_l through (15) on forest \mathcal{G}_l .
 Obtain $p^{k+1}|_{\mathcal{E}_l}$ by $K_l^\top p_l^{k+1} = v_l - f_l$.
end for
end for

and Sra [3] proposed a generalization of the taut-string approach to the case of weighted total variation which runs in $\mathcal{O}(n)$. The approach proposed by Johnson [31] also runs in $\mathcal{O}(n)$ time, works for weighted total variation and has good practical performance. Furthermore, a more memory efficient implementation of Johnson's algorithm generalization to trees was proposed by Kolmogorov et al. [34] – which appears to be state-of-the-art. We use our implementation of the algorithm proposed by Kolmogorov et al. [34] to find the exact minimizer on each tree. The algorithm computes derivatives of messages $\widehat{M}_i : \mathbb{R} \rightarrow \mathbb{R}$ and $M_{i,j} : \mathbb{R} \rightarrow \mathbb{R}$ for $(i, j) \in \mathcal{E}$ and $i, j \in \mathcal{V}$ in the order from leaves toward the root, which are defined as the following:

$$\widehat{M}_i(u) = \frac{1}{2}(u - f_{l,i})^2 + \sum_{(k,i) \in \mathcal{E}} M_{k,i}(u),$$

$$M_{i,j}(u) = \min_{u \in \mathbb{R}} \left[\widehat{M}_i(u) + \omega_{i,j}|u_j - u| \right].$$

The derivatives are denoted as $\widehat{m}_i := \widehat{M}'_i$, and $m_{i,j} := M'_{i,j}$. The procedure is summarized in Algorithm 1.

After running Algorithm 1 for each forest \mathcal{G}_l , the solution $p^{k+1}|_{\mathcal{E}_l}$ to (14) is given by the optimality condition

$$K_l^\top p^{k+1}|_{\mathcal{E}_l} = v_l - f_l.$$

Since each \mathcal{G}_l is a forest, the corresponding K_l^\top matrix has full column rank which implies that the linear system has a unique solution. Rows of K_l^\top corresponding to leaf nodes in the tree contain exactly one nonzero element. Therefore, we solve the linear equation by starting from leaves toward the root. Consistency on the branch nodes is guaranteed due to the uniqueness of the solution. Even though we only discussed the case of total variation $F(\cdot) = \|\cdot\|_1$, we remark that for various other choices of F (e.g., Huber penalty) efficient solvers on trees are conceivable.

name	Instance		None [11]		Diagonal [42]		Nested Forest		Linear Forest	
	$\frac{ V }{1024}$	$\frac{ \mathcal{E} }{ V }$	it	time[s]	it	time[s]	it	time[s]	it	time[s]
synthetic										
rmf-long.n2	64	2.87	–	–	–	–	1794	62.9 (+0.7)	5070	24.4 (+1.2)
rmf-wide.n2	32	2.84	–	–	–	–	159	2.6 (+0.3)	23518	62.2 (+0.6)
wash-rlg-long.n1024	64	2.99	73309	134.1 (+9.4)	14333	49.7 (+0.0)	5848	373.9 (+0.8)	18798	108.7 (+1.3)
bisection										
horse-48112	47	2.99	–	–	21593	38.3 (+0.0)	964	33.9 (+0.3)	19145	90.4 (+2.8)
alue7065-33338	33	1.61	–	–	23218	21.5 (+0.0)	2499	21.4 (+0.17)	49452	99.4 (+1.7)
stereo										
BVZ-venus1*	162	1.99	9068	20.8 (+7.9)	3741	20.2 (+0.0)	1111	90.7 (+0.9)	414	1.9 (+0.1)
BVZ-venus2*	162	1.99	10099	24.4 (+7.9)	3124	17.2 (+0.0)	1065	88.8 (+0.9)	384	1.7 (+0.1)
KZ2-sawtooth1	310	2.91	96974	736.5 (+31.8)	3468	51.7 (+0.0)	336	102.8 (+3.0)	526	14.9 (+6.1)
KZ2-sawtooth2	294	2.79	95849	675.3 (+33.1)	3520	48.2 (+0.0)	432	124.8 (+2.9)	652	15.7 (+5.8)
misc vision										
texture_graph	9	4.76	4860	1.6 (+1.45)	3554	1.9 (+0.0)	1091	9.12 (+0.1)	1669	1.7 (+0.3)
lazybrush-mangagirl*	579	1.99	–	–	–	–	13318	5727.3 (+3.3)	6330	95.4 (+0.4)
imgsegm-ferro	231	3.98	3594	32.2 (+75.9)	5806	95.0 (+0.0)	786	276.8 (+4.4)	775	25.6 (+3.1)

Table 2: We compare the number of iterations and time required to reach a relative primal dual gap of less than 10^{-10} on various graph cut instances. The time for constructing the preconditioners is shown in brackets (for “None” it is the time taken to estimate $\sigma_{\max}(K)$). “–” indicates that the method failed to reach the desired tolerance within 10^5 iterations. “*” indicates that the graph has grid topology.

4 Numerical Validation

The preconditioned PDHG algorithm (2)–(3) for total variation minimization on weighted graphs is summarized in Algorithm 2. We assume that the primal update can be efficiently computed, e.g. if G is separable.

For the experiments we compare the proposed preconditioners to the unpreconditioned variant of PDHG ($S = I$, $T = I$, $s = \|K\|$, $t = \|K\|$), the diagonal preconditioners from [42] with choice of $\alpha = 1$ and $s = t = 1$. When using the proposed preconditioners we employ the balanced step size choice $s = \sqrt{L}$, $t = \sqrt{L}$.

We implemented all algorithms in MATLAB, whereas time critical parts such as the total variation solver on a tree (Algorithm 1) were implemented in C++.

4.1 Generalized Fused Lasso

The fused lasso [48], also known as the Rudin-Osher-Fatemi (ROF) model [45] to the image processing community is readily generalized to graphs:

$$\min_{u \in \mathbb{R}^V} \frac{1}{2} \|u - f\|^2 + \|Ku\|_1. \quad (16)$$

Despite its simplicity, this model has a plethora of applications in statistics [52], machine learning [27, 8] and computer vision [47, 39], often as a subproblem in sequential convex programming for nonconvex minimization.

We solve (16) using the accelerated PDHG variant ([11, Algorithm 2], $\gamma = 0.25$) since the energy (16) is 1-strongly convex. In Table 1 we compare the number of iterations

required to solve (16) on small random graphs. We stop the algorithm once the relative primal-dual gap drops below 10^{-12} . It can be observed that there is a clear correlation between the condition number $\kappa(T^{-1/2}K)$ and the number of required iterations, validating the discussion from Section 2. The optimal preconditioning based on matroid partitioning performs best.

We further validate our preconditioner on the maximum flow benchmark [26]¹. It is well known (cf. [10, 42]) that the minimum cut in a flow network can be obtained by thresholding the minimizing argument of (16). We remark that graph cuts can be efficiently found by highly specialized combinatorial solvers such as that in [26]. The point of this experiment is, however, to compare different preconditioners for continuous first-order algorithms on challenging real-world graphs.

In Table 2, we show iterations and run time for the proposed forest preconditioners, the unpreconditioned PDHG algorithm and the diagonal preconditioners [42]. Due to the size of the graphs, matroid partitioning is intractable, and we resort to the greedy nested forest and linear forest approaches. Combinatorial preconditioning consistently leads to a significant decrease in iterations. In all except one case, the overall lowest run time is achieved either by linear forest decomposition or greedy nested forest decomposition. Despite the huge decrease in the total number of outer iterations for the nested forest preconditioning, in some cases, the overall run time is worse than without preconditioning. This motivates the construction of preconditioners like the greedy linear forests. These are clearly

¹<http://www.cs.tau.ac.il/~sagihed/ibfs/>



 Input (625×391) + Scribbles Result (11 PDHG iterations)

Figure 2: Interactive image segmentation. Combinatorial preconditioning works particularly well for graphs with underlying grid topology. The multiclass segmentation shown on the right is obtained after only 11 PDHG iterations.

suboptimal with respect to condition number, but yield a good balance between efficient resolution of the backward step and the number of outer iterations. The chains on regular grid achieve the best of both worlds and lead to an improvement in runtime of an order of magnitude.

4.2 Multiclass Segmentation

As a second example, we consider the multiclass total variation segmentation problem

$$\min_{u \in \mathbb{R}^{|\mathcal{V}| \times C}} \sum_{i \in \mathcal{V}} \delta\{u_i \in \Delta_C\} + \varepsilon H(u_i) + \langle u_i, \rho_i \rangle + \|Ku\|_1,$$

under entropic regularization $H(u_i) = \sum_{c=1}^C u_{i,c} \log u_{i,c}$. $\delta\{u_i \in \Delta_C\}$ is the indicator function of the C -dimensional unit simplex $\Delta_C \subset \mathbb{R}^C$.

The above model has various important applications in vision [53, 40] and transductive learning [22]. As the entropy term renders the objective $(1/\varepsilon)$ -strongly convex we can again use the accelerated PDHG algorithm ([11, Algorithm 2] with $\gamma = 1/\varepsilon$). We set $\varepsilon = 1$ for all experiments. Note that the model fits the general saddle point problem (1) under the choice $F(p) = \|p\|_1$, $G(u) = \sum_{i \in \mathcal{V}} \delta\{u_i \in \Delta_C\} + \varepsilon H(u_i) + \langle u_i, \rho_i \rangle$. To compute the proximal subproblem in (2) we first observe that $G^*(v) = \sum_{i \in \mathcal{V}} \varepsilon \log \left(\sum_{c=1}^C \exp((v_{i,c} - \rho_{i,c})/\varepsilon) \right)$.

Thanks to Moreau’s identity we reduce (2) to the proximal evaluation of G^* , for which a few Newton iterations suffice. In Table 3 we compare the performance of our combinatorial preconditioners on two of the aforementioned applications. For the transductive learning scenario, we follow the procedure described in [22] to generate a k -nearest neighbour graph ($k = 10$) on the synthetic “three moons” data set. As in [22], the data term $\rho \in \mathbb{R}^{|\mathcal{V}| \times C}$ specifies the correct labels for 5% of the points. We report a similar final accuracy (98.9%) as the authors of [22]. As seen in Table 3, the nested forest preconditioning performs best.

For the second application, we consider interactive image segmentation [40]. Following that paper, we compute the data term ρ from user scribbles (see Fig. 2). The weights

name	None [11] it (time[s])	Diag. [42] it (time[s])	Nest. Forest it (time[s])	Lin. Forest it (time[s])
3MOONS	333 (2.8)	474 (8.2)	88 (1.1)	303 (3.1)
icgbench-1*	113 (47.0)	131 (41.3)	52 (41.9)	13 (3.1)
icgbench-2*	124 (54.7)	159 (54.1)	58 (47.9)	11 (2.6)
icgbench-3*	78 (29.1)	95 (25.9)	42 (27.2)	9 (1.6)

Table 3: We compare iterations and time (in brackets) required for PDHG under various choices of preconditioner to reach a relative primal-dual gap less than $5 \cdot 10^{-4}$. On general graphs, the greedy nested forests perform well while on regular grids (indicated with “*”) the linear forest decomposition into chains works best.

are chosen based on the input image $I \in \mathbb{R}^{|\mathcal{V}| \times 3}$ as

$$\omega_{i,j} = \exp(-\xi \|I_i - I_j\|^2), \quad \forall (i,j) \in \mathcal{E}.$$

This essentially encourages the segmentation boundary to coincide with image discontinuities. We use a fixed scale parameter $\xi = 0.1$ in all experiments. The performance of the different preconditioning strategies is shown in Table 3. Due to the underlying grid topology the natural linear forest decomposition into chains can be employed, which outperforms the other preconditioners and PDHG without preconditioning by an order of magnitude. While the nested forest preconditioner is competitive w.r.t. iterations, the proximal subproblem on the large spanning tree is very expensive. In contrast, the subproblem on the short chains can be computed efficiently.

5 Conclusion

We proposed a novel combinatorial preconditioner for proximal algorithms on weighted graphs. The preconditioner is driven by a disjoint decomposition of the edge set into forests. Our theoretical analysis provides conditions under which such a decomposition achieves an optimal condition number. Furthermore, we have shown how provably optimal preconditioners can be obtained: on two-dimensional regular grids by a splitting into horizontal and vertical chains, on general graphs by means of matroid partitioning. We additionally proposed two fast heuristics to construct reasonable preconditioners on large scale graphs. We demonstrated how the resulting scaled proximal evaluations can be carried out by means of an efficient message passing algorithm on trees. In an extensive numerical evaluation we confirmed practical gains of preconditioning in terms of overall runtime as well as outer iteration numbers.

Acknowledgements

We thank Thomas Windheuser for fruitful discussions on combinatorial preconditioning. We gratefully acknowledge the support of the ERC Consolidator Grant 3D Reloaded.

References

- [1] Z. Allen-Zhu, Y. Li, R. Oliveira, and A. Wigderson. Much faster algorithms for matrix scaling. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2017.
- [2] K. J. Arrow, L. Hurwicz, and H. Uzawa. *Studies in Linear and Nonlinear Programming*. Stanford University Press, 1958.
- [3] A. Barbero and S. Sra. Modular proximal optimization for multidimensional total-variation regularization. *arXiv:1411.0589*, 2014.
- [4] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2:183–202, 2009.
- [5] S. Becker and M. J. Fadili. A quasi-Newton proximal splitting method. In *Proceedings of the 26th International Conference on Neural Information Processing Systems, NIPS*, 2012.
- [6] K. Benzi, V. Kalofolias, X. Bresson, and P. Vandergheynst. Song recommendation with non-negative matrix factorization and graph total variation. In *Proceedings of the 42nd IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2016.
- [7] K. Bredies and H. Sun. Preconditioned Douglas–Rachford splitting methods for convex-concave saddle-point problems. *SIAM J. Numer. Anal.*, 53:421–444, 2015.
- [8] X. Bresson, T. Laurent, D. Uminsky, and J. Von Brecht. Multiclass total variation clustering. In *Proceedings of the 27th International Conference on Neural Information Processing Systems, NIPS*, 2013.
- [9] A. E. Brouwer and W. H. Haemers. *Spectra of Graphs*. Springer, 2012.
- [10] A. Chambolle and J. Darbon. On total variation minimization and surface evolution using parametric maximum flows. *Int. J. Comput. Vis.*, 84:288–307, 2009.
- [11] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vis.*, 40:120–145, 2011.
- [12] A. Chambolle and T. Pock. On the ergodic convergence rates of a first-order primal–dual algorithm. *Math. Program.*, 159:253–287, 2016.
- [13] A. Chambolle and T. Pock. An introduction to continuous optimization for imaging. *Acta Numer.*, 25:161–319, 2016.
- [14] L. Condat. A direct algorithm for 1D total variation denoising. *IEEE Signal Process. Lett.*, 20:1054–1057, 2013.
- [15] P. L. Davies and A. Kovac. Local extremes, runs, strings and multiresolution. *Ann. Stat.*, 29:1–48, 2001.
- [16] S. Diamond and S. Boyd. Stochastic matrix-free equilibration. *J. Optim. Theory Appl.*, 172:436–454, 2017.
- [17] L. Dümbgen and A. Kovac. Extensions of smoothing via taut strings. *Electron. J. Stat.*, 3:41–75, 2009.
- [18] J. Edmonds. Minimum partition of a matroid into independent subsets. *J. Res. Natl. Bur. Standards*, 69B:67–72, 1965.
- [19] E. Esser, X. Zhang, and T. F. Chan. A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM J. Imaging Sci.*, 3:1015–1046, 2010.
- [20] C. Fougner and S. Boyd. Parameter selection and pre-conditioning for a graph form solver. *arXiv:1503.08366*, 2015.
- [21] M. P. Friedlander and G. Goh. Efficient evaluation of scaled proximal operators. *Electron. Trans. Numer. Anal.*, 46:1–22, 2017.
- [22] C. Garcia-Cardona, E. Merkurjev, A. L. Bertozzi, A. Flenner, and A. G. Percus. Multiclass data segmentation using diffuse interface methods on graphs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(8):1600–1613, 2014.
- [23] P. Giselsson and S. Boyd. Diagonal scaling in Douglas-Rachford splitting and ADMM. In *Proceedings of the 53rd IEEE Conference on Decision and Control, CDC*, 2014.
- [24] P. Giselsson and S. Boyd. Preconditioning in fast dual gradient methods. In *Proceedings of the 53rd IEEE Conference on Decision and Control, CDC*, 2014.
- [25] P. Giselsson and S. Boyd. Metric selection in fast dual forward–backward splitting. *Automatica*, 62:1–10, 2015.
- [26] A. Goldberg, S. Hed, H. Kaplan, R. Tarjan, and R. Werneck. Maximum flows by incremental breadth-first search. *European Symposium on Algorithms, ALGO ESA*, 2011.
- [27] M. Hein and S. Setzer. Beyond spectral clustering – tight relaxations of balanced graph cuts. In *Proceedings of the 25th International Conference on Neural Information Processing Systems, NIPS*, 2011.
- [28] M. Hein, S. Setzer, L. Jost, and S. S. Rangapuram. The total variation on hypergraphs – learning on hypergraphs revisited. In *Proceedings of the 27th International Conference on Neural Information Processing Systems, NIPS*, 2013.
- [29] A. J. Hoffman. On approximate solutions of systems of linear inequalities. *J. Res. Natl. Bur. Standards*, 49:263–265, 1952.

- [30] W. Imrich and I. Peterin. Recognizing Cartesian products in linear time. *Discrete Math.*, 307:472–483, 2007.
- [31] N. A. Johnson. A dynamic programming algorithm for the fused lasso and l_0 -segmentation. *J. Comput. Graph. Stat.*, 22(2):246–260, 2013.
- [32] A. Joshi. *Topics in Optimization and Sparse Linear Systems*. PhD thesis, UIUC, 1997.
- [33] D. E. Knuth. Matroid partitioning. Technical report, Stanford University, 1973.
- [34] V. Kolmogorov, T. Pock, and M. Rolinek. Total variation on a tree. *SIAM J. Imaging Sci.*, 9:605–636, 2016.
- [35] J. D. Lee, Y. Sun, and M. A. Saunders. Proximal Newton-type methods for minimizing composite functions. *SIAM J. Optim.*, 24:1420–1443, 2014.
- [36] C. S. A. Nash-Williams. Decomposition of finite graphs into forests. *J. London Math. Soc.*, 39:12–12, 1964.
- [37] I. Necoara, Y. Nesterov, and F. Glineur. Linear convergence of first order methods for non-strongly convex optimization. *arXiv:1504.06298*, 2015.
- [38] Y. Nesterov. A method for solving the convex programming problem with convergence rate $\mathcal{O}(1/k^2)$. *Soviet Mathematics Doklady*, 269:543–547, 1983.
- [39] R. A. Newcombe, S. J. Lovegrove, and A. J. Davison. DTAM: Dense tracking and mapping in real-time. In *Proceedings of the 13th International Conference on Computer Vision, ICCV*, 2011.
- [40] C. Nieuwenhuis and D. Cremers. Spatially varying color distributions for interactive multilabel segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(5):1234–1247, 2013.
- [41] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1:123–231, 2013.
- [42] T. Pock and A. Chambolle. Diagonal preconditioning for first order primal-dual algorithms in convex optimization. In *Proceedings of the 13th International Conference on Computer Vision, ICCV*, 2011.
- [43] T. Pock, D. Cremers, H. Bischof, and A. Chambolle. An algorithm for minimizing the piecewise smooth Mumford-Shah functional. In *Proceedings of the 11th International Conference on Computer Vision, ICCV*, 2009.
- [44] B. T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4:1–17, 1964.
- [45] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
- [46] D. A. Spielman. Algorithms, graph theory, and linear equations in Laplacian matrices. In *Proceedings of the International Congress of Mathematicians, ICM*, 2010.
- [47] J. Stühmer, S. Gumhold, and D. Cremers. Real-time dense geometry from a handheld camera. In *Proceedings of the 32nd DAGM Symposium on Pattern Recognition*, 2010.
- [48] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *J. R. Stat. Soc.*, 67(1):91–108, 2005.
- [49] P. M. Vaidya. Solving linear equations with symmetric diagonally dominant matrices by constructing good preconditioners. (A talk based on the manuscript was presented at the IMA Workshop on Graph Theory and Sparse Matrix Computation), 1991.
- [50] P.-W. Wang and C.-J. Lin. Iteration complexity of feasible descent methods for convex optimization. *J. Mach. Learn. Res.*, 15:1523–1548, 2014.
- [51] H. Weyl. Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). *Mathematische Annalen*, 71:441–479, 1912.
- [52] B. Xin, Y. Kawahara, Y. Wang, and W. Gao. Efficient generalized fused Lasso and its application to the diagnosis of Alzheimer’s disease. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, 2014.
- [53] C. Zach, D. Gallup, J.-M. Frahm, and M. Niethammer. Fast global labeling for real-time stereo using multiple plane sweeps. In *Proceedings of the Vision, Modeling and Visualization Workshop*, 2008.
- [54] K. Zhong, I. E. H. Yen, I. S. Dhillon, and P. Ravikumar. Proximal quasi-Newton for computationally intensive ℓ_1 -regularized M -estimators. In *Proceedings of the 28th International Conference on Neural Information Processing Systems, NIPS*, 2014.
- [55] M. Zhu and T. F. Chan. An efficient primal-dual hybrid gradient algorithm for total variation image restoration. CAM Reports 08-34, UCLA, 2008.