

Efficient Linear Programming for Dense CRFs

Thalaiyasingam Ajanthan¹, Alban Desmaison², Rudy Bunel², Mathieu Salzmann³, Philip H.S. Torr²,
and M. Pawan Kumar^{2,4}

¹Australian National University & Data61, CSIRO

²Department of Engineering Science, University of Oxford

³Computer Vision Laboratory, EPFL

⁴Alan Turing Institute

Abstract

The fully connected conditional random field (CRF) with Gaussian pairwise potentials has proven popular and effective for multi-class semantic segmentation. While the energy of a dense CRF can be minimized accurately using a linear programming (LP) relaxation, the state-of-the-art algorithm is too slow to be useful in practice. To alleviate this deficiency, we introduce an efficient LP minimization algorithm for dense CRFs. To this end, we develop a proximal minimization framework, where the dual of each proximal problem is optimized via block coordinate descent. We show that each block of variables can be efficiently optimized. Specifically, for one block, the problem decomposes into significantly smaller subproblems, each of which is defined over a single pixel. For the other block, the problem is optimized via conditional gradient descent. This has two advantages: 1) the conditional gradient can be computed in a time linear in the number of pixels and labels; and 2) the optimal step size can be computed analytically. Our experiments on standard datasets provide compelling evidence that our approach outperforms all existing baselines including the previous LP based approach for dense CRFs.

1. Introduction

In the past few years, the dense conditional random field (CRF) with Gaussian pairwise potentials has become popular for multi-class image-based semantic segmentation. At the origin of this popularity lies the use of an efficient filtering method [1], which was shown to lead to a linear time mean-field inference strategy [12]. Recently, this filtering method was exploited to minimize the dense CRF energy using other, typically more effective, continuous relaxation methods [6]. Among the relaxations considered in [6], the linear programming (LP) relaxation provides strong theoretical guarantees on the quality of the solution [9, 15].

In [6], the LP was minimized via projected subgradient descent. While relying on the filtering method, computing the subgradient was shown to be *linearithmic* in the number of pixels, but not *linear*. Moreover, even with the use of a line search strategy, the algorithm required a large number of iterations to converge, making it inefficient.

We introduce an iterative LP minimization algorithm for

a dense CRF with Gaussian pairwise potentials which has *linear* time complexity per iteration. To this end, instead of relying on a standard subgradient technique, we propose to make use of the proximal method [20]. The resulting proximal problem has a smooth dual, which can be efficiently optimized using block coordinate descent. We show that each block of variables can be optimized efficiently. Specifically, for one block, the problem decomposes into significantly smaller subproblems, each of which is defined over a single pixel. For the other block, the problem can be optimized via the Frank-Wolfe algorithm [8, 16]. We show that the conditional gradient required by this algorithm can be computed efficiently. In particular, we modify the filtering method of [1] such that the conditional gradient can be computed in a time *linear* in the number of pixels and labels. Besides this linear complexity, our approach has two additional benefits. First, it can be initialized with the solution of a faster, less accurate algorithm, such as mean-field [12] or the difference of convex (DC) relaxation of [6], thus speeding up convergence. Second, the optimal step size of our iterative procedure can be obtained analytically, thus preventing the need to rely on an expensive line search procedure.

We demonstrate the effectiveness of our algorithm on the MSRC and Pascal VOC 2010 [7] segmentation datasets. The experiments evidence that our algorithm is significantly faster than the state-of-the-art LP minimization technique of [6]. Furthermore, it yields assignments whose energies are much lower than those obtained by the baselines [6, 12]. Altogether, our framework constitutes the first efficient and effective minimization algorithm for dense CRFs with Gaussian pairwise potentials. Our code is available at <https://github.com/oval-group/DenseCRF> and a more detailed version of the paper can be found at <https://arxiv.org/abs/1611.09718>.

2. Preliminaries

Before introducing our method, let us first provide some background on the dense CRF model and its LP relaxation.

Dense CRF energy function. A dense CRF is defined on a set of n random variables $\mathcal{X} = \{X_1, \dots, X_n\}$, where each random variable X_a takes a label $x_a \in \mathcal{L}$, with $|\mathcal{L}| = m$. For a given labelling \mathbf{x} , the energy associated with a

pairwise dense CRF can be expressed as

$$E(\mathbf{x}) = \sum_{a=1}^n \phi_a(x_a) + \sum_{a=1}^n \sum_{b=1, b \neq a}^n \psi_{ab}(x_a, x_b), \quad (1)$$

where ϕ_a and ψ_{ab} denote the *unary potentials* and *pairwise potentials*, respectively. The unary potentials define the data cost and the pairwise potentials the smoothness cost.

Gaussian pairwise potentials. Similarly to [6, 12], we consider Gaussian pairwise potentials, which have the following form:

$$\psi_{ab}(x_a, x_b) = \mu(x_a, x_b) \sum_c w^{(c)} k(\mathbf{f}_a^{(c)}, \mathbf{f}_b^{(c)}), \quad (2)$$

$$k(\mathbf{f}_a, \mathbf{f}_b) = \exp\left(\frac{-\|\mathbf{f}_a - \mathbf{f}_b\|^2}{2}\right).$$

Here, $\mu(x_a, x_b)$ is referred to as the *label compatibility function* and the mixture of Gaussian kernels as the *pixel compatibility function*. The weights $w^{(c)} \geq 0$ define the mixture coefficients, and $\mathbf{f}_a^{(c)} \in \mathbb{R}^{d^{(c)}}$ encodes features associated to the random variable X_a , where $d^{(c)}$ is the feature dimension. For semantic segmentation, each pixel in an image corresponds to a random variable. In practice, as in [6, 12], we then use the position and RGB values of a pixel as features, and assume the label compatibility function to be the Potts model, that is, $\mu(x_a, x_b) = \mathbb{1}[x_a \neq x_b]$. These potentials have proven useful to obtain fine grained labellings in segmentation tasks [12].

Integer programming formulation. An alternative way of representing a labelling is by defining indicator variables $y_{a:i} \in \{0, 1\}$, where $y_{a:i} = 1$ if and only if $x_a = i$. Using this notation, the energy minimization problem can be written as the following Integer Program (IP):

$$\begin{aligned} \min_{\mathbf{y}} E(\mathbf{y}) &= \sum_a \sum_i \phi_{a:i} y_{a:i} + \sum_{a, b \neq a} \sum_{i, j} \psi_{ab:ij} y_{a:i} y_{b:j}, \\ \text{s.t.} \quad \sum_i y_{a:i} &= 1 \quad \forall a \in \{1 \dots n\}, \\ y_{a:i} &\in \{0, 1\} \quad \forall a \in \{1 \dots n\}, \quad \forall i \in \mathcal{L}. \end{aligned} \quad (3)$$

Here, we use the shorthand $\phi_{a:i} = \phi_a(i)$ and $\psi_{ab:ij} = \psi_{ab}(i, j)$. The first set of constraints ensure that each random variable is assigned exactly one label. Note that the value of the objective function is equal to the energy of the labelling encoded by \mathbf{y} .

Linear programming relaxation. By relaxing the binary constraints of the indicator variables in (3) and using the fact that the label compatibility function is the Potts model, the linear programming relaxation [9] of (3) is defined as

$$\begin{aligned} \min_{\mathbf{y}} \tilde{E}(\mathbf{y}) &= \sum_a \sum_i \phi_{a:i} y_{a:i} + \sum_{a, b \neq a} \sum_i K_{ab} \frac{|y_{a:i} - y_{b:i}|}{2}, \\ \text{s.t.} \quad \mathbf{y} \in \mathcal{M} &= \left\{ \mathbf{y} \mid \begin{array}{l} \sum_i y_{a:i} = 1, a \in \{1 \dots n\} \\ y_{a:i} \geq 0, a \in \{1 \dots n\}, i \in \mathcal{L} \end{array} \right\}, \end{aligned} \quad (4)$$

where $K_{ab} = \sum_c w^{(c)} k(\mathbf{f}_a^{(c)}, \mathbf{f}_b^{(c)})$. For integer labellings, the LP objective $\tilde{E}(\mathbf{y})$ has the same value as the

IP objective $E(\mathbf{y})$. The above relaxation is the same as the standard LP relaxation [4] for the Potts model and it provides an integrality gap of 2. The result in [17] means that it is unlikely (unless the Unique Games Conjecture is false) that a better relaxation can be designed for this problem. Using standard solvers to minimize this LP would require the introduction of $\mathcal{O}(n^2)$ variables (see Appendix A.1), making it intractable. Therefore the non-smooth objective of Eq. (4) has to be optimized directly. This was handled using projected subgradient descent in [6], which also turns out to be inefficient in practice. In this paper, we introduce an efficient algorithm to tackle this problem while maintaining *linear* scaling in both space and time complexity.

3. Proximal minimization for LP relaxation

Our goal is to design an efficient minimization strategy for the LP relaxation in (4). To this end, we propose to use the proximal minimization algorithm [20]. This guarantees monotonic decrease in the objective value, enabling us to leverage faster, less accurate methods for initialization. Furthermore, the additional quadratic regularization term makes the dual problem smooth, enabling the use of more sophisticated optimization methods. In the remainder of this paper, we detail this approach and show that each iteration has linear time complexity. In practice, our algorithm converges in a small number of iterations, thereby making the overall approach computationally efficient.

The proximal minimization algorithm [20] is an iterative method that, given the current estimate of the solution \mathbf{y}^k , solves the problem

$$\begin{aligned} \min_{\mathbf{y}} \quad & \tilde{E}(\mathbf{y}) + \frac{1}{2\lambda} \|\mathbf{y} - \mathbf{y}^k\|^2, \\ \text{s.t.} \quad & \mathbf{y} \in \mathcal{M}, \end{aligned} \quad (5)$$

where λ sets the strength of the proximal term.

Note that (5) consists of piecewise linear terms and a quadratic regularization term. Specifically, the piecewise linear term comes from the pairwise term $|y_{a:i} - y_{b:i}|$ in (4) that can be reformulated as $\max\{y_{a:i} - y_{b:i}, y_{b:i} - y_{a:i}\}$. The proximal term $\|\mathbf{y} - \mathbf{y}^k\|^2$ provides the quadratic regularization. In this section, we introduce a new algorithm that is tailored to this problem. In particular, we optimally solve the Lagrange dual of (5) in a block-wise fashion.

3.1. Dual formulation

The dual problem has three variables, namely, $\alpha = \{\alpha_{ab:i}^1, \alpha_{ab:i}^2 \mid a, b \neq a, i \in \mathcal{L}\}$, $\beta = \{\beta_a \mid a \in \{1 \dots n\}\}$ and $\gamma = \{\gamma_{a:i}, a \in \{1 \dots n\}, i \in \mathcal{L}\}$. Here, we introduce two matrices that will be useful to write the dual problem compactly. Specifically, the matrices $A \in \mathbb{R}^{nm \times p}$ (where $p = 2n(n-1)m$) and $B \in \mathbb{R}^{nm \times n}$ are defined such that

$$\begin{aligned} (A\alpha)_{a:i} &= - \sum_{b \neq a} (\alpha_{ab:i}^1 - \alpha_{ab:i}^2 + \alpha_{ba:i}^2 - \alpha_{ba:i}^1), \\ (B\beta)_{a:i} &= \beta_a. \end{aligned} \quad (6)$$

Algorithm 1 Proximal minimization of LP (PROX-LP)

Require: Initial solution $\mathbf{y}^0 \in \mathcal{M}$ and the dual objective g

for $k \leftarrow 0 \dots K$ **do**

$$A\boldsymbol{\alpha}^0 \leftarrow \mathbf{0}, \quad \boldsymbol{\beta}^0 \leftarrow \mathbf{0}, \quad \boldsymbol{\gamma}^0 \leftarrow \mathbf{0}$$

▷ Feasible initialization

for $t \leftarrow 0 \dots T$ **do**

$$(\boldsymbol{\beta}^t, \boldsymbol{\gamma}^t) \leftarrow \underset{\boldsymbol{\beta}, \boldsymbol{\gamma}}{\operatorname{argmin}} g(\boldsymbol{\alpha}^t, \boldsymbol{\beta}, \boldsymbol{\gamma})$$

▷ Sec. 3.2.1

$$\tilde{\mathbf{y}}^t \leftarrow \lambda (A\boldsymbol{\alpha}^t + B\boldsymbol{\beta}^t + \boldsymbol{\gamma}^t - \boldsymbol{\phi}) + \mathbf{y}^k$$

▷ Current primal solution, may be infeasible

$$A\mathbf{s}^t \leftarrow \text{conditional gradient of } g, \text{ computed using } \tilde{\mathbf{y}}^t$$

▷ Sec. 3.2.2

$$\delta \leftarrow \text{optimal step size given } (\mathbf{s}^t, \boldsymbol{\alpha}^t, \tilde{\mathbf{y}}^t)$$

▷ Sec. 3.2.2

$$A\boldsymbol{\alpha}^{t+1} \leftarrow (1 - \delta)A\boldsymbol{\alpha}^t + \delta A\mathbf{s}^t$$

▷ Frank-Wolfe update on $\boldsymbol{\alpha}$

$$\mathbf{y}^{k+1} \leftarrow P_{\mathcal{M}}(\tilde{\mathbf{y}}^t)$$

▷ Project the primal solution to the feasible set \mathcal{M}

This lets us write the Lagrange dual of (5) as

$$\min_{\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}} g(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}) = \frac{\lambda}{2} \|A\boldsymbol{\alpha} + B\boldsymbol{\beta} + \boldsymbol{\gamma} - \boldsymbol{\phi}\|^2 + \langle A\boldsymbol{\alpha} + B\boldsymbol{\beta} + \boldsymbol{\gamma} - \boldsymbol{\phi}, \mathbf{y}^k \rangle - \langle \mathbf{1}, \boldsymbol{\beta} \rangle, \quad (7)$$

$$\text{s.t. } \gamma_{a:i} \geq 0 \quad \forall a \in \{1 \dots n\} \quad \forall i \in \mathcal{L},$$

$$\boldsymbol{\alpha} \in \mathcal{C} = \left\{ \boldsymbol{\alpha} \mid \begin{array}{l} \alpha_{ab:i}^1 + \alpha_{ab:i}^2 = \frac{K_{ab}}{2}, \quad a, b \neq a, \quad i \in \mathcal{L} \\ \alpha_{ab:i}^1, \alpha_{ab:i}^2 \geq 0, \quad a, b \neq a, \quad i \in \mathcal{L} \end{array} \right\}.$$

Given the dual variables, the corresponding primal variables can be obtained using the KKT conditions [3] as

$$\mathbf{y} = \lambda (A\boldsymbol{\alpha} + B\boldsymbol{\beta} + \boldsymbol{\gamma} - \boldsymbol{\phi}) + \mathbf{y}^k. \quad (8)$$

We refer the reader to Appendix A.1 in the supplementary material for the details of this derivation.

3.2. Algorithm

The dual problem (7), in its standard form, can only be tackled using projected gradient descent. However, by separating the variables based on the type of the feasible domains, we propose an efficient block coordinate descent approach. Each of these blocks are amenable to more sophisticated optimization, resulting in a computationally efficient algorithm. As the dual problem is strictly convex and smooth, the optimal solution is still guaranteed. For $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$, the problem decomposes over the pixels, as shown in 3.2.1, therefore making it efficient. The minimization with respect to $\boldsymbol{\alpha}$ is over a compact domain, which can be efficiently tackled using the Frank-Wolfe algorithm [8, 16]. Our complete algorithm is summarized in Algorithm 1. In the following sections, we discuss each step in more detail.

3.2.1 Optimizing over $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$

We first turn to the problem of optimizing over $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ while $\boldsymbol{\alpha}^t$ is fixed. Since the dual variable $\boldsymbol{\beta}$ is unconstrained, the minimum value of the dual objective g is attained when $\nabla_{\boldsymbol{\beta}} g(\boldsymbol{\alpha}^t, \boldsymbol{\beta}, \boldsymbol{\gamma}) = 0$. Differentiating with respect to $\boldsymbol{\beta}$ and setting the derivatives to zero yields

$$\boldsymbol{\beta} = B^T (A\boldsymbol{\alpha}^t + \boldsymbol{\gamma} - \boldsymbol{\phi}) / m. \quad (9)$$

Note that, now, $\boldsymbol{\beta}$ is a function of $\boldsymbol{\gamma}$. We therefore substitute $\boldsymbol{\beta}$ in (7) and minimize over $\boldsymbol{\gamma}$. Interestingly, the resulting problem can be optimized independently for each pixel, with each subproblem being an m dimensional quadratic program (QP) with nonnegativity constraints, where m is the number of labels. For a pixel a , this QP has the form

$$\min_{\boldsymbol{\gamma}_a} \frac{1}{2} \boldsymbol{\gamma}_a^T Q \boldsymbol{\gamma}_a + \langle \boldsymbol{\gamma}_a, Q ((A\boldsymbol{\alpha}^t)_a - \boldsymbol{\phi}_a) + \mathbf{y}_a^k \rangle, \quad (10)$$

$$\text{s.t. } \boldsymbol{\gamma}_a \geq \mathbf{0}.$$

Here, $\boldsymbol{\gamma}_a$ denotes the vector $\{\gamma_{a:i} \mid i \in \mathcal{L}\}$ and $Q = \lambda (I - \mathbf{1}/m) \in \mathbb{R}^{m \times m}$, with I the identity matrix and $\mathbf{1}$ a matrix of all ones.

We use the algorithm of [27] to efficiently optimize every such QP. In our case, due to the structure of the matrix Q , the time complexity of an iteration is linear in the number of labels. Hence, the overall time complexity of optimizing over $\boldsymbol{\gamma}$ is $\mathcal{O}(nm)$. Once the optimal $\boldsymbol{\gamma}$ is computed for a given $\boldsymbol{\alpha}^t$, the corresponding optimal $\boldsymbol{\beta}$ is given by Eq. (9). More details are provided in Appendix A.2.

3.2.2 Optimizing over $\boldsymbol{\alpha}$

We now turn to the problem of optimizing over $\boldsymbol{\alpha}$ given $\boldsymbol{\beta}^t$ and $\boldsymbol{\gamma}^t$. To this end, we use the Frank-Wolfe algorithm [8], which has the advantage of being projection free. Furthermore, for our specific problem, we show that the required conditional gradient can be computed efficiently and the optimal step size can be obtained analytically.

Conditional gradient computation. The conditional gradient with respect to $\boldsymbol{\alpha}$ is obtained by solving the linearization problem

$$\mathbf{s} = \underset{\mathbf{s} \in \mathcal{C}}{\operatorname{argmin}} \langle \mathbf{s}, \nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^t, \boldsymbol{\beta}^t, \boldsymbol{\gamma}^t) \rangle. \quad (11)$$

Here, $\nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^t, \boldsymbol{\beta}^t, \boldsymbol{\gamma}^t)$ denotes the gradient of the dual objective function with respect to $\boldsymbol{\alpha}$ evaluated at $(\boldsymbol{\alpha}^t, \boldsymbol{\beta}^t, \boldsymbol{\gamma}^t)$.

Importantly, we show that the conditional gradient has an analytical form, given by

$$(A\mathbf{s})_{a:i} = - \sum_b (K_{ab} \mathbb{1}[\tilde{y}_{a:i}^t \geq \tilde{y}_{b:i}^t] - K_{ab} \mathbb{1}[\tilde{y}_{a:i}^t \leq \tilde{y}_{b:i}^t]), \quad (12)$$

where $\tilde{\mathbf{y}}^t$ is the current (infeasible) primal solution computed using Eq. (8). We refer the reader to Appendix A.3 for the detailed derivation.

Note that Eq. (12) has the same form as the LP subgradient (Eq. (20) in [6]). This is not a surprising result. In fact, it has been shown that, for certain problems, there exists a duality relationship between subgradients and conditional gradients [2]. To compute this subgradient, the state-of-the-art algorithm proposed in [6] has a time complexity linearithmic in the number of pixels. Unfortunately, since this constitutes a critical step of both our algorithm and that of [6], such a linearithmic cost greatly affects their efficiency. In Section 4, however, we show that this complexity can be reduced to linear, thus effectively leading to a speedup of an order of magnitude in practice.

Optimal step size. One of the main difficulties of using an iterative algorithm, whether subgradient or conditional gradient descent, is that its performance depends critically on the choice of the step size. Here, we can analytically compute the optimal step size that results in the maximum decrease in the objective for the given descent direction. As shown in Appendix A.4, the resulting step is given by

$$\delta = P_{[0,1]} \left(\frac{\langle A\boldsymbol{\alpha}^t - A\mathbf{s}^t, \tilde{\mathbf{y}}^t \rangle}{\lambda \|A\boldsymbol{\alpha}^t - A\mathbf{s}^t\|^2} \right). \quad (13)$$

Here, $P_{[0,1]}$ denotes the projection to the interval $[0, 1]$, that is, clipping the value to lie in $[0, 1]$.

Memory efficiency. For a dense CRF, the dual variable $\boldsymbol{\alpha}$ requires $\mathcal{O}(n^2m)$ storage, which becomes infeasible since n is the number of pixels in an image. Note, however, that $\boldsymbol{\alpha}$ always appears in the product $\tilde{\boldsymbol{\alpha}} = A\boldsymbol{\alpha}$ in Algorithm 1. Therefore, we only store the variable $\tilde{\boldsymbol{\alpha}}$, which reduces the storage complexity to $\mathcal{O}(nm)$.

3.2.3 Summary

To summarize, our method has the following desirable qualities of an efficient iterative algorithm. First, it can benefit from an initial solution obtained by a faster but less accurate algorithm, such as mean-field or DC relaxation. Second, with our choice of a quadratic proximal term, the dual of the proximal problem can be efficiently optimized in a block-wise fashion. Specifically, the dual variables $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are computed efficiently by minimizing one small QP (of dimension the number of labels) for each pixel independently. The remaining dual variable $\boldsymbol{\alpha}$ is optimized using the Frank-Wolfe algorithm, where the conditional gradient is computed in linear time, and the optimal step size is obtained analytically. Overall, the time complexity of one iteration of our algorithm is $\mathcal{O}(nm)$. To the best of our knowledge, this constitutes the first LP minimization algorithm for dense CRFs that has linear time iterations. We denote this algorithm as PROX-LP.

4. Fast conditional gradient computation

The algorithm described in the previous section assumes that the conditional gradient (Eq. (12)) can be computed efficiently. Note that Eq. (12) contains two terms that are similar up to sign and order of the label constraint in the indicator function. To simplify the discussion, let us focus on the first term and on a particular label i , which we will not explicitly write in the remainder of this section. The second term in Eq. (12) and the other labels can be handled in the same manner. With these simplifications, we need to efficiently compute an expression of the form

$$\forall a \in \{1 \dots n\}, \quad v'_a = \sum_b k(\mathbf{f}_a, \mathbf{f}_b) \mathbb{1}[y_a \geq y_b], \quad (14)$$

with $y_a, y_b \in [0, 1]$ and $\mathbf{f}_a, \mathbf{f}_b \in \mathbb{R}^d$ for all $a, b \in \{1 \dots n\}$.

The usual way of speeding up computations involving such Gaussian kernels is by using the efficient filtering method [1]. This approximate method has proven accurate enough for similar applications [6, 12]. In our case, due to the ordering constraint $\mathbb{1}[y_a \geq y_b]$, the symmetry is broken and the direct application of the filtering method is impossible. In [6], the authors tackled this problem using a divide-and-conquer strategy, which lead to a time complexity of $\mathcal{O}(d^2n \log(n))$. In practice, this remains a prohibitively high run time, particularly since gradient computations are performed many times over the course of the algorithm. Here, we introduce a more efficient method.

Specifically, we show that the term in Eq. (14) can be computed in $\mathcal{O}(Hdn)$ time (where H is a small constant defined in Section 4.2), at the cost of additional storage. In practice, this leads to a speedup of one order of magnitude. Below, we first briefly review the original filtering algorithm and then explain our modified algorithm that efficiently handles the ordering constraints.

4.1. Original filtering method

In this section, we assume that the reader is familiar with the permutohedral lattice based filtering method [1]. Due to space constraint, only a brief overview is provided. We refer the interested reader to the original paper [1].

In [1], each pixel $a \in \{1 \dots n\}$ is associated with a tuple (\mathbf{f}_a, v_a) , which we call a *feature point*. The elements of this tuple are the feature $\mathbf{f}_a \in \mathbb{R}^d$ and the value $v_a \in \mathbb{R}$. Note that, in our case, $v_a = 1$ for all pixels. At the beginning of the algorithm, the feature points are embedded in a d -dimensional hyperplane tessellated by the *permutohedral lattice* (the hexagon shape shown in Fig. 1). The vertices of this permutohedral lattice are called *lattice points*.

Once the permutohedral lattice is constructed, the algorithm performs three main steps: *splatting*, *blurring* and *slicing*. During splatting, for each lattice point, the values of the neighbouring feature points are accumulated using barycentric interpolation. Next, during blurring, the values of the lattice points are convolved with a one dimensional truncated Gaussian kernel along each feature dimen-

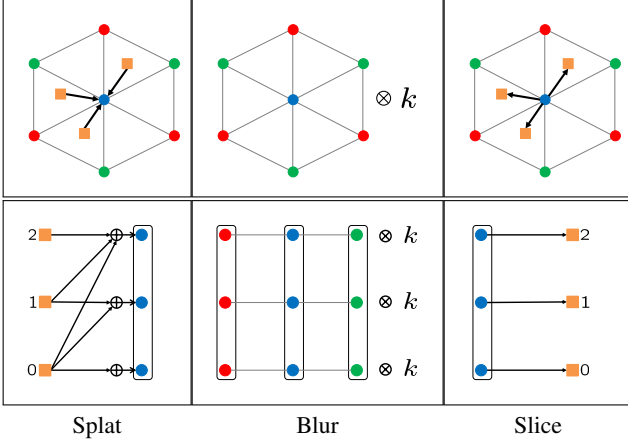


Figure 1: *The hexagon made of triangles represents the permutohedral lattice with $d = 2$, where the feature points are denoted with squares and the lattice points with circles. **Top row:** Original filtering method. The barycentric interpolation is denoted by an arrow and k here is the truncated Gaussian kernel. See para-3 in Sec. 4.1. **Bottom row:** Our modified filtering method. Here, $H = 3$, and the figure therefore illustrates 3 lattices. We write the bin number of each feature point next to it. See para-2 in Sec. 4.2.*

sion separately. Finally, during slicing, the resulting values of the lattice points are propagated back to the feature points using the same barycentric weights. These steps are explained graphically in the top row of Fig. 1. The pseudocode of the algorithm is given in Appendix B.1. The time complexity of this algorithm is $\mathcal{O}(dn)$ [1, 12], and the complexity of the permutohedral lattice creation $\mathcal{O}(d^2n)$. Since the approach in [6] creates multiple lattices at every iteration, the overall complexity of this approach is $\mathcal{O}(d^2n \log(n))$.

Note that, in this original algorithm, there is no notion of score y_a associated with each pixel. In particular, during splatting, the values v_a are accumulated to the neighbouring lattice points without considering their scores. Therefore, this algorithm cannot be directly applied to handle our ordering constraint $\mathbb{1}[y_a \geq y_b]$.

4.2. Modified filtering method

We now introduce a filtering-based algorithm that can handle ordering constraints. To this end, we uniformly discretize the continuous interval $[0, 1]$ into H different discrete bins, or levels. Note that each pixel, or feature point, belongs to exactly one of these bins, according to its corresponding score. We then propose to instantiate H permutohedral lattices, one for each level $h \in \{0 \dots H - 1\}$.

To handle the ordering constraints, we then modify the splatting step in the following manner. A feature point belonging to bin q is splat to the permutohedral lattices corresponding to levels $q \leq h < H$. Blurring is then performed independently in each individual permutohedral lattice. This guarantees that a feature point will only influ-

ence the values of the feature points that belong to the same level or higher ones. In other words, a feature point b influences the value of a feature point a only if $y_a \geq y_b$. Finally, during the slicing step, the value of a feature point belonging to level q is recovered from the q^{th} permutohedral lattice. Our algorithm is depicted graphically in the bottom row of Fig. 1. Its pseudocode is provided in Appendix B.2. Note that, while discussed for constraints of the form $\mathbb{1}[y_a \geq y_b]$, this algorithm can easily be adapted to handle $\mathbb{1}[y_a \leq y_b]$ constraints, which are required for the second term in Eq. (12).

Overall, our modified filtering method has a time complexity of $\mathcal{O}(Hdn)$ and a space complexity of $\mathcal{O}(Hdn)$. Note that the complexity of the lattice creation is still $\mathcal{O}(d^2n)$ and can be reused for each of the H instances. Moreover, as opposed to the method in [6], this operation is performed only once, during the initialization step. In practice, we were able to choose H as small as 10, thus achieving a substantial speedup compared to the divide-and-conquer strategy of [6]. By discretizing the interval $[0, 1]$, we add another level of approximation to the overall algorithm. However, this approximation can be eliminated by using a dynamic data structure (see Appendix B.2.1).

5. Related work

We review the past work on three different aspects of our work in order to highlight our contributions.

Dense CRF. The fully-connected CRF has become increasingly popular for semantic segmentation. It is particularly effective at preventing oversmoothing, thus providing better accuracy at the boundaries of objects. As a matter of fact, in a complementary direction, many methods have now proposed to combine dense CRFs with convolutional neural networks [5, 22, 28] to achieve state-of-the-art performance on segmentation benchmarks.

The main challenge that had previously prevented the use of dense CRFs is their computational cost at inference, which, naively, is $\mathcal{O}(n^2)$ per iteration. In the case of Gaussian pairwise potential, the efficient filtering method of [1] proved to be key to the tractability of inference in the dense CRF. While an approximate method, the accuracy of the computation proved sufficient for practical purposes. This was first observed in [12] for the specific case of mean-field inference. More recently, several continuous relaxations, such as QP, DC and LP, were also shown to be applicable to minimizing the dense CRF energy by exploiting this filtering procedure in various ways [6]. Unfortunately, while tractable, minimizing the LP relaxation, which is known to provide the best approximation to the original labelling problem, remained too slow in practice [6]. Our algorithm is faster both theoretically and empirically. Furthermore, and as evidenced by our experiments, it yields lower energy values than any existing dense CRF inference strategy.

LP relaxation. There are two ways to relax the integer program (3) to a linear program, depending on the label compatibility function: 1) the standard LP relaxation [4]; and 2) the LP relaxation specialized to the Potts model [9]. There are many notable works on minimizing the standard LP relaxation on sparse CRFs. This includes the algorithms that directly make use the dual of this LP [10, 11, 25] and those based on a proximal minimization framework [18, 21]. Unfortunately, all of the above algorithms exploit the sparsity of the problem, and they would yield an $\mathcal{O}(n^2)$ cost per iteration in the fully-connected case. In this work, we focus on the Potts model based LP relaxation for dense CRFs and provide an algorithm whose iterations have time complexity $\mathcal{O}(n)$. Even though we focus on the Potts model, as pointed out in [6], this LP relaxation can be extended to general label compatibility functions using a hierarchical Potts model [14].

Frank-Wolfe. The optimization problem of structural support vector machines (SVM) has a form similar to our proximal problem. The Frank-Wolfe algorithm [8] was shown to provide an effective and efficient solution to such a problem via block-coordinate optimization [16]. Several works have recently focused on improving the performance of this algorithm [19, 23] and extended its application domain [13]. Our work draws inspiration from this structural SVM literature, and makes use of the Frank-Wolfe algorithm to solve a subtask of our overall LP minimization method. Efficiency, however, could only be achieved thanks to our modification of the efficient filtering procedure to handle ordering constraints.

To the best of our knowledge, our approach constitutes the first LP minimization algorithm for dense CRFs to have linear time iterations. Our experiments demonstrate the importance of this result on both speed and labelling quality. Being fast, our algorithm can be incorporated in any end-to-end learning framework, such as [28]. We therefore believe that it will have a significant impact on future semantic segmentation results, and potentially in other applications.

6. Experiments

In this section, we first discuss two variants that further speedup our algorithm and some implementation details. We then turn to the empirical results.

6.1. Accelerated variants

Empirically we observed that, our algorithm can be accelerated by restricting the optimization procedure to affect only relevant subsets of labels and pixels. These subsets can be identified from an intermediate solution of PROX-LP. In particular, we remove the label i from the optimization if $y_{a:i} < 0.01$ for all pixels a . In other words, the score of a label i is insignificant for all the pixels. We denote this version as PROX-LP $_{\ell}$. Similarly, we optimize over a pixel only if it is *uncertain* in choosing a label. Here, a pixel a is called

uncertain if $\max_i y_{a:i} < 0.95$. In other words, no label has a score higher than 0.95. The intuition behind this strategy is that, after a few iterations of PROX-LP $_{\ell}$, most of the pixels are labelled correctly, and we only need to fine tune the few remaining ones. In practice, we limit this restricted set to 10% of the total number of pixels. We denote this accelerated algorithm as PROX-LP $_{\text{acc}}$. As shown in our experiments, PROX-LP $_{\text{acc}}$ yields a significant speedup at virtually no loss in the quality of the results.

6.2. Implementation details

In practice, we initialize our algorithm with the solution of the best continuous relaxation algorithm, which is called DC $_{\text{neg}}$ in [6]. The parameters of our algorithm, such as the proximal regularization constant λ and the stopping criteria, are chosen manually. A small value of λ leads to easier minimization of the proximal problem, but also yields smaller steps at each proximal iteration. We found $\lambda = 0.1$ to work well in all our experiments. We fixed the maximum number of proximal steps (K in Algorithm 1) to 10, and each proximal step is optimized for a maximum of 5 Frank-Wolfe iterations (T in Algorithm 1). In all our experiments the number of levels H is fixed to 10.

6.3. Segmentation results

We evaluated our algorithm on the MSRC and Pascal VOC 2010 [7] segmentation datasets, and compare it against mean-field inference (MF) [12], the best performing continuous relaxation method of [6] (DC $_{\text{neg}}$) and the subgradient based LP minimization method of [6] (SG-LP). Note that, in [6], the LP was initialized with the DC $_{\text{neg}}$ solution and optimized for 5 iterations. Furthermore, the LP optimization was performed on a subset of labels identified by the DC $_{\text{neg}}$ solution in a similar manner to the one discussed in Section 6.1. We refer to this algorithm as SG-LP $_{\ell}$. For all the baselines, we employed the respective authors' implementations that were obtained from the web or through personal communication. Furthermore, for all the algorithms, the integral labelling is computed from the fractional solution using the *argmax* rounding scheme.

For both datasets, we used the same splits and unary potentials as in [12]. The pairwise potentials were defined using two kernels: a spatial kernel and a bilateral one [12]. For each method, the kernel parameters were cross validated on validation data using Spearmint [24] (with a budget of 2 days). To be able to compare energy values, we then evaluated all methods with the same parameters. In other words, for each dataset, each method was run several times with different parameter values. Note that, on MSRC, cross-validation was performed on the less accurate ground truth provided with the original dataset. Nevertheless, we evaluated all methods on the accurate ground truth annotations provided by [12].

The results for the parameters tuned for DC $_{\text{neg}}$ on the MSRC and Pascal datasets are given in Table 1. Here MF5




Figure 2: Assignment energy as a function of time with the parameters tuned for DC_{neg} for an image in (left) MSRC and (right) Pascal. A zoomed-in version is shown next to each plot. Except MF, all other algorithms are initialized with DC_{neg} . Note that PROX-LP clearly outperforms $SG-LP_\ell$ by obtaining much lower energies in fewer iterations. Furthermore, the accelerated versions of our algorithm obtain roughly the same energy as PROX-LP but significantly faster.

	Algorithm	MF5	MF	DC_{neg}	SG- LP_ℓ	PROX- LP	PROX- LP_ℓ	PROX- LP_{acc}	Avg. E ($\times 10^3$)	Avg. T (s)	Acc.	IoU
MSRC	MF5	-	0	0	0	0	0	0	8078.0	0.2	79.33	52.30
	MF	96	-	0	0	0	0	0	8062.4	0.5	79.35	52.32
	DC_{neg}	96	96	-	0	0	0	0	3539.6	1.3	83.01	57.92
	$SG-LP_\ell$	96	96	90	-	3	1	1	3335.6	13.6	83.15	58.09
	PROX-LP	96	96	94	92	-	13	45	1274.4	23.5	83.99	59.66
	$PROX-LP_\ell$	96	96	95	94	81	-	61	1189.8	6.3	83.94	59.50
	$PROX-LP_{acc}$	96	96	95	94	49	31	-	1340.0	3.7	84.16	59.65
Pascal	MF5	-	13	0	0	0	0	0	1220.8	0.8	79.13	27.53
	MF	2	-	0	0	0	0	0	1220.8	0.7	79.13	27.53
	DC_{neg}	99	99	-	-	0	0	0	629.5	3.7	80.43	28.60
	$SG-LP_\ell$	99	99	95	-	5	12	12	617.1	84.4	80.49	28.68
	PROX-LP	99	99	95	84	-	32	50	507.7	106.7	80.63	28.53
	$PROX-LP_\ell$	99	99	86	86	64	-	43	502.1	22.1	80.65	28.29
	$PROX-LP_{acc}$	99	99	86	86	46	39	-	507.7	14.7	80.58	28.45

Table 1: Results on the MSRC and Pascal datasets with the parameters tuned for DC_{neg} . We show: the percentage of images where the row method strictly outperforms the column one on the final integral energy, the average integral energy over the test set, the average run time, the segmentation accuracy and the intersection over union score. Note that all versions of our algorithm obtain much lower energies than the baselines. Interestingly, while our fully accelerated version does slightly worse in terms of energy, it is the best in terms of the segmentation accuracy in MSRC.

denotes the mean-field algorithm run for 5 iterations. In Fig. 2, we show the assignment energy as a function of time for an image in MSRC (the tree image in Fig. 3) and for an image in Pascal (the sheep image in Fig. 3). Furthermore, we provide some of the segmentation results in Fig. 3.

In summary, $PROX-LP_\ell$ obtains the lowest integral energy in both datasets. Furthermore, our fully accelerated version is the fastest LP minimization algorithm and always outperforms the baselines by a great margin in terms of energy. From Fig. 3, we can see that $PROX-LP_{acc}$ marks most of the crucial pixels (e.g., object boundaries) as *uncertain*, and optimizes over them efficiently and effectively. Note that, on top of being fast, $PROX-LP_{acc}$ obtains the highest accuracy in MSRC for the parameters tuned for DC_{neg} .

To ensure consistent behaviour across different energy

parameters, we ran the same experiments for the parameters tuned for MF. In this setting, all versions of our algorithm again yield significantly lower energies than the baselines. Due to space constraint, results on MSRC for selected algorithms are summarized in Table 2, and complete results for this parameter setting are given in Appendix C.2.

Our observation that a lower energy does not necessarily result in improved segmentation is an important one (similar behaviour was observed in [6, 26]). Indeed, this lack of correlation between the dense CRF energy and the segmentation accuracy highlights the importance of performing a more thorough analysis of the dense CRF model.

6.4. Modified filtering method

We then compare our modified filtering method, described in Section 4, with the divide-and-conquer strategy




Figure 3: *Qualitative results with the parameters tuned for DC_{neg} for an image in (top) MSRC and (bottom) Pascal. The uncertain pixels identified by $PROX-LP_{acc}$ are marked in white. Note that, in both images, all versions of our algorithm obtain visually good segmentations. In addition, $PROX-LP_{acc}$ identifies the crucial pixels (object boundaries) as uncertain and efficiently optimizes over them. Furthermore, in the MSRC image, the improvement of $PROX-LP_{acc}$ over the baselines is clearly visible, and the final segmentation is virtually the same as the accurate ground truth. (Best viewed in color)*

Algorithm	Avg. E ($\times 10^3$)	Avg. T (s)	Acc.	IoU
MF	1053.6	13.0	83.86	59.75
DC_{neg}	812.7	2.8	83.50	59.67
$PROX-LP_{acc}$	295.9	7.9	83.03	58.97

Table 2: *Results on MSRC with the parameters tuned for MF. Our method obtains the lowest energy, however, as expected, since the parameters are tuned for MF, this inference strategy yields the best accuracy.*

of [6]. To this end, we evaluated both algorithms on one of the Pascal VOC test images (the sheep image in Fig. 3), but varying the image size, the number of labels and the Gaussian kernel standard deviation. Note that, to generate a plot for one variable, the other variables are fixed to their respective standard values. The standard value for the number of pixels is 187500, for the number of labels 21, and for the standard deviation 1. For this experiment, the conditional gradients were computed from a random primal solution \tilde{y}^t . In Fig. 4, we show the speedup of our modified filtering approach over the one of [6] as a function of the number of pixels and labels. As shown in Appendix C.4, the speedup with respect to the kernel standard deviation is roughly constant. The timings were averaged over 10 runs, and we observed only negligible timing variations between the different runs.

In summary, our modified filtering method is 10 – 65 times faster than the state-of-the-art algorithm of [6]. Furthermore, note that all versions of our algorithm operate in the region where the speedup is around 45 – 65.

7. Discussion

We have introduced the first LP minimization algorithm for dense CRFs with Gaussian pairwise potentials whose iterations are linear in the number of pixels and labels. Thanks to the efficiency of our algorithm and to the tightness of the LP relaxation, our approach yields much lower

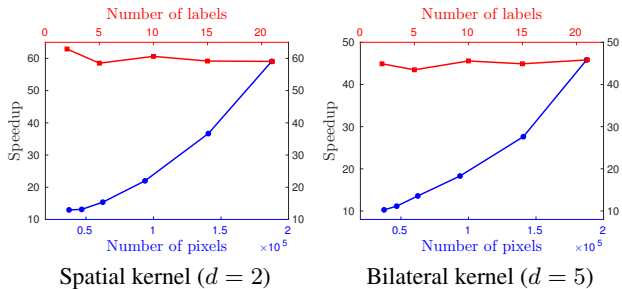


Figure 4: *Speedup of our modified filtering algorithm over the divide-and-conquer strategy of [6] on a Pascal image. Note that our speedup grows with the number of pixels and is approximately constant with respect to the number of labels. (Best viewed in color)*

energy values than state-of-the-art dense CRF inference methods. Furthermore, our experiments demonstrated that, with the right set of energy parameters, highly accurate segmentation results can be obtained with our algorithm. The speed and effective energy minimization of our algorithm make it a perfect candidate to be incorporated in an end-to-end learning framework, such as [28]. This, we believe, will be key to further improving the accuracy of deep semantic segmentation architectures.

8. Acknowledgements

This work was supported by the EPSRC, the ERC grant ERC- 2012-AdG 321162-HELIOS, the EPSRC/MURI grant ref EP/N019474/1, the EPSRC grant EP/M013774/1, the EPSRC programme grant Seebibyte EP/M013774/1, the Microsoft Research PhD Scholarship, ANU PhD Scholarship and Data61 Scholarship. Data61 (formerly NICTA) is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

References

- [1] A. Adams, J. Baek, and M. Davis. Fast high-dimensional filtering using the permutohedral lattice. *Computer Graphics Forum*, 2010. 1, 4, 5
- [2] F. Bach. Duality between subgradient and conditional gradient methods. *SIAM Journal on Optimization*, 2015. 4
- [3] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2009. 3
- [4] C. Chekuri, S. Khanna, J. Naor, and L. Zosin. A linear programming formulation and approximation algorithms for the metric labeling problem. *SIAM Journal on Discrete Mathematics*, 2004. 2, 6
- [5] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. Yuille. Semantic image segmentation with deep convolutional nets and fully connected CRFs. *ICLR*, 2014. 5
- [6] A. Desmaison, R. Bunel, P. Kohli, P. Torr, and P. Kumar. Efficient continuous relaxations for dense CRF. *ECCV*, 2016. 1, 2, 4, 5, 6, 7, 8
- [7] M. Everingham, L. Van Gool, C. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (VOC) challenge. *IJCV*, 2010. 1, 6
- [8] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 1956. 1, 3, 6
- [9] J. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relationships: metric labeling and markov random fields. *Journal of the ACM*, 2002. 1, 2, 6
- [10] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *PAMI*, 2006. 6
- [11] N. Komodakis, N. Paragios, and G. Tziritas. MRF energy minimization and beyond via dual decomposition. *PAMI*, 2011. 6
- [12] P. Krähenbühl and V. Koltun. Efficient inference in fully connected CRFs with gaussian edge potentials. *NIPS*, 2011. 1, 2, 4, 5, 6
- [13] R. Krishnan, S. Lacoste-Julien, and D. Sontag. Barrier Frank-Wolfe for marginal inference. *NIPS*, 2015. 6
- [14] P. Kumar and D. Koller. MAP estimation of semi-metric MRFs via hierarchical graph cuts. *UAI*, 2009. 6
- [15] P. Kumar, V. Kolmogorov, and P. Torr. An analysis of convex relaxations for MAP estimation of discrete MRFs. *JMLR*, 2009. 1
- [16] S. Lacoste-Julien, M. Jaggi, M. Schmidt, and P. Pletscher. Block-coordinate Frank-Wolfe optimization for structural SVMs. *ICML*, 2012. 1, 3, 6
- [17] R. Manokaran, J. Naor, P. Raghavendra, and R. Schwartz. SDP Gaps and UGC Hardness for Multiway Cut, 0-Extension and Metric Labeling. *STOC*, 2008. 2
- [18] O. Meshi, M. Mahdavi, and A. Schwing. Smooth and strong: MAP inference with linear convergence. *NIPS*, 2015. 6
- [19] A. Osokin, J. Alayrac, I. Lukasewitz, P. Dokania, and S. Lacoste-Julien. Minding the gaps for block Frank-Wolfe optimization of structured SVMs. *ICML*, 2016. 6
- [20] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 2014. 1, 2
- [21] P. Ravikumar, A. Agarwal, and M. Wainwright. Message-passing for graph-structured linear programs: proximal projections, convergence and rounding schemes. *ICML*, 2008. 6
- [22] A. Schwing and R. Urtasun. Fully connected deep structured networks. *CoRR*, 2015. 5
- [23] N. Shah, V. Kolmogorov, and C. Lampert. A multi-plane block-coordinate Frank-Wolfe algorithm for training structural SVMs with a costly max-oracle. *CVPR*, 2015. 6
- [24] J. Snoek, H. Larochelle, and R. Adams. Practical bayesian optimization of machine learning algorithms. *NIPS*, 2012. 6
- [25] M. Wainwright, T. Jaakkola, and A. Willsky. MAP estimation via agreement on trees: message-passing and linear programming. *Information Theory*, 2005. 6
- [26] P. Wang, C. Shen, and A. van den Hengel. Efficient SDP inference for fully-connected CRFs based on low-rank decomposition. *CVPR*, 2015. 7
- [27] X. Xiao and D. Chen. Multiplicative iteration for nonnegative quadratic programming. *Numerical Linear Algebra with Applications*, 2014. 3
- [28] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. Torr. Conditional random fields as recurrent neural networks. *ICCV*, 2015. 5, 6, 8