

A fast solver for truncated-convex priors: quantized-convex split moves^{*}

Anna Jeziarska¹, Hugues Talbot¹, Olga Veksler², and Daniel Wesierski³

¹ Université Paris-Est, Laboratoire d'Informatique Gaspard-Monge, France,
`{first.last}@univ-paris-est.fr`

² University of Western Ontario, Canada, `olga@csd.uwo.ca`

³ Telecom SudParis, France, `daniel.wesierski@it-sudparis.eu`

Abstract. This paper addresses the problem of minimizing multilabel energies with truncated convex priors. Such priors are known to be useful but difficult and slow to optimize because they are not convex. We propose two novel classes of binary Graph-Cuts (GC) moves, namely the *convex move* and the *quantized move*. The moves are complementary. To significantly improve efficiency, the label range is divided into even intervals. The quantized move tends to efficiently put pixel labels into the correct intervals for the energy with truncated convex prior. Then the convex move assigns the labels more precisely within these intervals for the same energy. The quantized move is a modified α -expansion move, adapted to handle a generalized Potts prior, which assigns a constant penalty to arguments above some threshold. Our convex move is a GC representation of the efficient Murota's algorithm. We assume that the data terms are convex, since this is a requirement for Murota's algorithm. We introduce *Quantized-Convex Split Moves* algorithm which minimizes energies with truncated priors by alternating both moves. This algorithm is a fast solver for labeling problems with a high number of labels and convex data terms. We illustrate its performance on image restoration.

Key words: Graph cuts, image restoration, non-convex prior, Potts model, Murota's algorithm.

1 Introduction

We consider the well-known combinatorial optimization problem defined as follows. Let $\mathcal{G}(\mathcal{V}, \mathcal{E})$ be an undirected graph with a set of edges \mathcal{E} and a set of vertices \mathcal{V} . The goal of our optimization problem is to restore an unknown x^* based on observations \bar{x} , under the condition that x takes values over a finite set of labels \mathcal{L} , representing e.g. grey level values in an image. Here we define \mathcal{L} as an ordered discrete set of labels $\{0, 1, \dots, L\}$ and x_u as the label assigned to node $u \in \mathcal{V}$. The unknown x^* is a minimum argument of the energy function:

$$E(x) = \sum_{u \in \mathcal{V}} \mathcal{D}(x_u) + \lambda \sum_{(u,v) \in \mathcal{E}} \mathcal{R}(x_u, x_v), \quad (1)$$

^{*} This work was supported by the Agence Nationale de la Recherche under grant ANR-09-EMER-004-03.

where λ is a positive real value. $D(x_u)$ is often called the data fidelity term and $\mathcal{R}(x_u, x_v)$ the regularization or smoothness term. A common choice of data term \mathcal{D} is a pixelwise distance $\mathcal{D} = |x_u - \bar{x}_u|^p$ between the desired labeling x and a reference \bar{x} , representing noisy acquired data, where p is a small positive integer, e.g. 1 or 2.

Many choices of \mathcal{R} lead to useful algorithms and results. A common model is the so-called Potts model, where $\mathcal{R}(x_u, x_v) = w_{uv} \min(1, |x_u - x_v|)$, and w_{uv} are spatially variant positive pairwise weights. This model corresponds to a piecewise constant prior. Other choices for \mathcal{R} include $\mathcal{R} = w_{uv}|x_u - x_v|^q$, where q is typically 1 or 2 for linear and quadratic priors respectively. The latter represents an “everywhere smooth” prior with good denoising properties and lack of staircase effect in the result, but with blurred boundaries. Better preservation of boundaries can be achieved with regularization term $\mathcal{R} = w_{uv} \min(T^q, |x_u - x_v|^q)$, where for $q = 1$ or $q = 2$ it is respectively called truncated linear or truncated quadratic [1]. More generally, a pairwise truncated convex prior can be formulated as:

$$\mathcal{R}(x_v, x_u) = \begin{cases} f(x_u - x_v) & \text{if } |x_u - x_v| < T \\ f(T) & \text{if } |x_u - x_v| \geq T \end{cases} \quad (2)$$

where f is a convex function with $f(0) = 0$. Discrete random field models characterized by such a prior are well known and extensively discussed in the literature. Their popularity in low level vision is due to their ability to capture natural image statistics [2]. Indeed, Nikolova [3] shows that the robustness of regularization terms depends on their characteristics at $\pm\infty$, and their differentiability at zero. Non-differentiable terms at zero reconstruct sharp edges well but lead to undesirable staircase effects. As a result, for the case of image restoration problems in the pixel domain, truncated regularization terms are more robust. In this way, truncated models may combine noise suppression with edge preservation. In general, depending on the application, a sharp (e.g truncated linear) or smooth (e.g. truncated quadratic) term might be desirable.

In the following, we introduce new GC algorithm solving optimization problem characterized by energy (1) and prior (2). In recent years, energy-based optimization methods using GC have become very popular in computer vision applications [4–6]. GC optimization has, for example, been applied to stereo-vision [7], multiview reconstruction [8], motion analysis [9], segmentation [10] and image restoration [11]. GC methods tend to provide optimal or near-optimal solutions to classical Markov Random Fields (MRF) problems, with some guarantees and in reasonable time, unlike earlier methods like Simulated Annealing (SA) [12] or Iterated Conditional Modes (ICM) [13]. From the algorithmic point of view, GC problems can be solved exactly when the energy is submodular, which was shown for the binary case (binary \mathcal{L}) in [14, 15] and for multilabel case in [16]. When energy E is not submodular, some GC methods can still be used, for instance the *move* algorithms [6, 17–19]

GC move algorithms have typically good theoretical guarantees for quality for certain sets of regularization terms containing truncated convex functions

considered in this paper. Classical move algorithms include expansion and swap moves [6]. More recently, improved moves have been proposed e.g. range moves and fusion moves [17, 18, 20, 19]. All are geared towards improving the quality of the solution and the speed of the algorithm. The time complexity of move algorithms usually increases steeply with the number of labels. For example, the worst-case complexity of swap moves is quadratic in the number of labels while range-moves perform even poorer. However, for problems where the number of labels is relatively low, these methods can be fast enough. Hence, move algorithms scale well with connectivity, are flexible with respect to data fidelity terms, but do not scale well with the number of labels.

It is worth noting that when \mathcal{R} is convex, e.g. in the non-truncated linear or quadratic cases, the energy E of (1) may be optimized exactly and efficiently [5, 11]. Moreover, Szeliski et al. [21] have shown that expansion and swap moves work well in practice for the Potts model. Conversely, in the truncated linear or quadratic cases, due to non-convexity and non-differentiability (at the truncation and also at zero for truncated linear regularization term), such optimization problems remain challenging. In the multilabel case, e.g. when the set of labels \mathcal{L} is not binary, the minimization problem of (1) is NP-hard.

The GC algorithms dedicated for energies with truncated convex priors e.g [17, 18, 22] have been developed to meet this challenge. We discuss them in detail in section 2. This group of algorithms can be extended with our *Quantized-Convex Split Moves*. This two-step approach produces results comparable to the current state-of-the-art move based algorithms, and yet outperforms them by a large factor in terms of time efficiency, especially when the number of labels is large. As these convex priors and the Potts model can be optimized efficiently with move methods, we split the label set into two parts, a regular quantized one that we optimize using a modified Potts model, and a remainder part, which we optimize using a convex framework. We propose two types of moves, which are complementary, namely the *convex move* and the *quantized move*. Our quantized move is a modified α -expansion move, adapted to cope with a generalized Potts prior taking zero value for arguments in the range $(-T, +T)$. Thus, it tends to efficiently put pixel labels into the right intervals. These approximate results are corrected by the convex move, which performs finer changes with respect to the previously chosen label. A new, more precise label is found within previously chosen interval. The convex move is a GC representation of an efficient Murota gradient descent algorithm [14, 23].

The rest of the paper is organized as follows. The description of our method in the context of mostly related work is given in section 2. We present our Quantized and Convex moves in section 3, and the Quantized-convex split moves algorithm in section 4. Then we provide experimental comparison of the different energy minimization methods in section 5, and conclude with section 6.

2 Related work

In recent years, many algorithms utilizing truncated regularization terms have been proposed. Apart from GC move algorithms, the sequential tree reweighted message passing (TRW-S) [24] has currently the most accurate results and provides a Lagrangian approximation of the dual energy, e.g. estimates the gap between current and globally optimal energies. However, it is relatively slow [25] and is not well suited to highly-connected graphs [26]. Belief propagation (BP) [27] methods, though fast, are not guaranteed to converge. GC methods were shown to outperform BP in several cases examined in [21]. Energies with truncated linear priors (truncated ℓ_1) may be optimized e.g. using α -expansions [6] or Gupta and Tardos [28] algorithm. The latter offers good theoretical properties, but it is not practical.

Veksler proposed in [17] to minimize energies with truncated convex priors by splitting the problem into several subproblems that are all convex with respect to the prior. Each subproblem is defined for subsets of pixels $\hat{u}, \hat{v} \in \mathcal{V}$ with labels $x_{\hat{u}} \in \mathcal{T}$ such that $\mathcal{T} \subset \mathcal{L}$ and $|x_{\hat{u}} - x_{\hat{v}}| \leq T$. Note that there exist many $\mathcal{T} \subset \mathcal{L}$ satisfying conditions $|x_{\hat{u}} - x_{\hat{v}}| \leq T$. Moreover, assuming that labels in $\mathcal{T} = \{\dots, t_{i-1}, t_i, t_{i+1}, \dots\}$ form a convex cone defined as $t_i + 1 = t_{i+1}$, one can assign T different $\mathcal{T} \subset \mathcal{L}$ to each x_u . According to the theorem presented in [17], the original energy with labeling \mathcal{L} is minimized with each subenergy having sublabeing \mathcal{T} . An algorithm that takes advantage of this property is the *range move*. Range move solves different subproblems for different choices of \mathcal{T} iteratively using an Ishikawa-like approach [5]. In this article, we show that using what we call a *convex move* instead of the Ishikawa approach, it is possible to consider all possible choices of $\mathcal{T} \subset \mathcal{L}$ such that $d_{\mathcal{T}} = T - 1$ simultaneously, where $d_{\mathcal{T}} = \max\{|x_{\hat{u}} - x_{\hat{v}}|, \{x_{\hat{u}}, x_{\hat{v}}\} \in \mathcal{T}\}$. This allows us to improve the time efficiency of the overall algorithm considerably. Additionally, we propose a *quantized move*, allowing for changes of x_u between $\mathcal{T}_1 \subset \mathcal{L}$ and $\mathcal{T}_2 \subset \mathcal{L}$ such that $\mathcal{T}_1 \cap \mathcal{T}_2 = \emptyset$. This further improves the time efficiency of our algorithm upon the range move. The proposed algorithm alternates iteratively between quantized and convex moves. Note that if $\mathcal{T}_1 \cap \mathcal{T}_2 = \emptyset$, the energy (1) is no longer convex with respect to the prior term. The advantage of the Ishikawa approach is that it guarantees a global minimum even with a non-convex data fidelity term, provided the prior is convex. This is particularly important for stereo vision. For the convex move introduced in this paper, the energy is guaranteed to decrease but the optimal solution is not secured.

More recent work by Kumar and Torr [18] is better grounded theoretically than Veksler's range move. The quality of the solution is guaranteed by bounds on the converged energy for truncated ℓ_1 and ℓ_2 , which are calculated with respect to $d_{\mathcal{T}}$, and equal $2 + \sqrt{2}$ if $d_{\mathcal{T}} = 2\sqrt{T}$ and $O(\sqrt{T})$ if $d_{\mathcal{T}} = \sqrt{T}$, for truncated ℓ_1 and ℓ_2 , respectively. However, according to the results presented in [18] the practical performance of both algorithms is similar for truncated ℓ_2 prior, although the greatest improvement is achieved for the truncated ℓ_1 prior. In terms of time efficiency, range move outperforms the approach proposed by Kumar and Torr, but not significantly. Similarly to the range move, authors use

the graph construction proposed by Ishikawa, but they introduce small modifications. Namely, they adopt the Ishikawa approach to deal with non-convex priors at the cost of not representing the energy exactly. Here we will not analyze our algorithm as a function of $d_{\mathcal{T}}$. The convex move in our quantized-convex split moves algorithm is associated with two sets: (1) the set of all possible $\mathcal{T} \subset \mathcal{L}$ with $d_{\mathcal{T}} = T - 1$ and (2) the set of all possible $\mathcal{T} \subset \mathcal{L}$ with $d_{\mathcal{T}} = T$.

In [22], authors proposed a hierarchical approach. The original problem was replaced by a series of r -HST metric labeling subproblems and obtained solutions were combined with α -expansion algorithm. The previously presented approximation bounds were improved. They are equal to $O(\ln(L))$ and $O((\gamma \ln(L))^2)$, $\gamma \geq 1$ for truncated ℓ_1 and ℓ_2 , respectively. However, this approach is computationally expensive.

3 Move algorithms

Move algorithms have been developed to solve multilabeling problems. According to the definition given in [17], a move algorithm is an iterative algorithm where $x_{n+1} \in M(x_n)$ and $M(x)$ is a “moves” space of x . The local minimum with respect to a set of moves is at x if $E(x') > E(x)$ for any $x' \in M(x)$. Each move algorithm is characterized by its space of “moves” $M(x)$.

In this section we describe two moves that we develop. The quantized move is closely related to α -expansion and convex move to Murota’s gradient descent algorithm. In section 4, we explain why linking these moves together leads to improvement of efficiency in the context of minimization of energy functions with truncated convex prior.

3.1 Quantized move

The main idea behind the quantized move is to divide the label range into equal subintervals of length T and, ideally, put pixel labels into the correct intervals, thus reducing the number of categories from the original range L to L/T . This greatly accelerates the execution time of the algorithm.

The proposed move algorithm minimizes the energy E_p with an arbitrary data fidelity term \mathcal{D}_p and a pairwise term defined as:

$$\mathcal{R}_p(x_v, x_u) = \begin{cases} 0 & \text{if } |x_u - x_v| < T \\ f(T) & \text{if } |x_u - x_v| \geq T, \end{cases} \quad (3)$$

where T is a positive integer value. This prior is potentially interesting for other applications, but here we will use it as an intermediate step for minimizing truncated convex priors.

A *quantized move* is a new labeling where x_u is either left as x_u or moved to a new value according to the following transformation:

$$\alpha(x_u, k) = \begin{cases} t_1^k & \text{if } x_u \leq t_1^k \\ t_T^k & \text{if } x_u \geq t_T^k, \end{cases} \quad (4)$$

where k is an integer belonging to a regular quantization of the label set \mathcal{L} , i.e.; $k \in \mathcal{K} = \{k_0, k_1, \dots, k_K\}$ such that $k_0 = 0$, $k_i = iT$, $i \in \mathbb{N}_+$, $KT \geq L$ and $(K-1)T < L$. Recall that L is the maximum label in \mathcal{L} . $\mathcal{T}^k = \{t_1^k, \dots, t_T^k\}$ is an ordered label set, such that $t_{i+1}^k = t_i^k + 1$. The values in \mathcal{T}^k change from $k - \frac{T}{2} + 1$ to $k + \frac{T}{2}$ and from $k - \frac{T}{2} + \frac{1}{2}$ to $k + \frac{T}{2} - \frac{1}{2}$ for odd and even T , respectively. The t_1^k and the t_T^k is a first and last element of set \mathcal{T}^k , respectively. The acceptable moves for a label depending on its current position are illustrated in Fig. 1.




Fig. 1. (a,b,c) illustrate the label moves when its current value is below, above and inside the considered interval \mathcal{T}^k (denoted by square brackets), respectively.

The set of quantized moves $M_Q(x)$ is then defined as the collection of moves for all $k \in \mathcal{K}$. Quantized moves act much like expansion moves in the case of a Potts model on a quantized subset of labels. We now prove that quantized moves are graph-representable and can be optimized by GC.

Proposition 1 *For the energy in (1) with a regularization term given by (3), the optimal quantized move (i.e. giving the maximum decrease in energy) can be computed with a graph cut.*

Proof: We show that quantized move satisfies all conditions specified in [15]. Let $b = \{b_u, \forall u \in \mathcal{V}\}$ be a binary vector coding a quantized move. Then the move can be described by a transformation function $\mathcal{B}(x^{(n)}, b)$ returning a new labeling $x^{(n+1)}$, based on b and $x^{(n)}$. Here (n) is the iteration number. The transformation function $\mathcal{B}_q(x^{(n)}, b)$ for a quantized move is given by:

$$x_u^{(n+1)} = \mathcal{B}_q(x_u^{(n)}, b_u) = \begin{cases} \alpha(x_u^{(n)}, k) & \text{if } b_u = 1 \\ x_u^{(n)} & \text{if } b_u = 0 \end{cases} \quad (5)$$

The considered move finds $b^* = \text{Argmin}_b E(\mathcal{B}_q(x^{(n)}, b))$, where $E(\mathcal{B}_q(x^{(n)}, b))$ is a pseudo-boolean energy, defined as $\sum_{u \in \mathcal{V}} \mathcal{D}(\mathcal{B}_q(x_u^{(n)}, b_u)) + \sum_{(u,v) \in \mathcal{E}} \mathcal{R}(\mathcal{B}_q(x_u^{(n)}, b_u), \mathcal{B}_q(x_v^{(n)}, b_v))$. Let us denote the pairwise term of the binary quantized move energy by $B(b_u, b_v)$, omitting $x^{(n)}$ from the notation for simplification. Then:

$$B(b_u, b_v) = \begin{cases} \mathcal{R}_p(x_u^{(n)}, x_v^{(n)}) & \text{if } b_u = 0, b_v = 0 \\ \mathcal{R}_p(x_u^{(n)}, \alpha(x_v^{(n)}, k)) & \text{if } b_u = 0, b_v = 1 \\ \mathcal{R}_p(\alpha(x_u^{(n)}, k), x_v^{(n)}) & \text{if } b_u = 1, b_v = 0 \\ \mathcal{R}_p(\alpha(x_u^{(n)}, k), \alpha(x_v^{(n)}, k)) & \text{if } b_u = 1, b_v = 1. \end{cases} \quad (6)$$

The pairwise term B needs to be submodular i.e.: $B(0,0) + B(1,1) \leq B(1,0) + B(0,1)$. Since for all n and k we have that $\mathcal{R}_p(\alpha(x_u^{(n)}, k), \alpha(x_v^{(n)}, k)) = 0$, the submodularity inequality takes the form:

$$\mathcal{R}_p(x_u^{(n)}, \alpha(x_v^{(n)}, k)) + \mathcal{R}_p(\alpha(x_u^{(n)}, k), x_v^{(n)}) \geq \mathcal{R}_p(x_u^{(n)}, x_v^{(n)}), \quad (7)$$

or equivalently

$$B(0,1) + B(1,0) \geq B(0,0). \quad (8)$$

The only case when $B(0,0)$ is not 0 is when neighbors x_u and x_v are at least T apart, i.e. $|x_u - x_v| \geq T$, in which case we have $B(0,0) = f(T)$. However, in this case either $B(0,1)$ or $B(1,0)$ or both are equal to $f(T)$, so the inequality is verified.

The problem of minimizing energy $E(\mathcal{B}_q(x^{(n)}, b))$ can be solved globally with respect to b using discrete maxflow-mincut methods [29]. Note that when $T = 1$ our quantized move reduces to the α -expansion move.

3.2 Convex moves

In the previous section, we showed how to assign the pixel values into the correct intervals, and now we propose a convex algorithm to optimize these values within these intervals. To achieve this, we view the steepest descent algorithm of Murota [14, 23] as a special case of GC move. The primal and a primal-dual algorithms proposed in [30] are also related to Murota's approach. Their convergence properties in the case of L^1 -convex functions have been proved. However, the case of non-convex data fidelity was not examined. This limitation can be viewed as disadvantage compared to Ishikawa approach [5], which guaranties a global minimum even for non-convex data fidelity. In contrast, both primal and primal-dual algorithms are more memory and time efficient than the non-iterative Ishikawa's method. We note that a function E_c with a non-convex data fidelity term and a convex pairwise regularization term $\mathcal{R}_c(x_u, x_v) = f(|x_u - x_v|)$ can still be minimized with the GC primal algorithm of [30]. The convex move is conceptually similar to the jump move [1]. However, the jump move processes pixels with odd and even values differently. As a consequence, Potts functions can be represented on jump-move graphs, whereas convex functions generally cannot.

As in the previous case (section 3.1), a convex move is described by a binary vector b and the transformation function $\mathcal{B}_c(x^{(n)}, b)$ defined as:

$$x_u^{(n+1)} = \mathcal{B}_c(x_u^{(n)}, b_u) = \begin{cases} x_u^{(n)} + s & \text{if } b_u = 1 \\ x_u^{(n)} & \text{if } b_u = 0, \end{cases} \quad (9)$$

where $s \in \mathcal{S}$ and \mathcal{S} is a set of discrete values from \mathbb{Z} . The convex move space $M_C(x)$ is then defined as the collection of convex moves for all $s \in \mathcal{S}$. We call the algorithm finding $b^* = \text{Argmin}_b E(\mathcal{B}_c(x^{(n)}, b))$ the *convex move algorithm*.

The pseudo-boolean prior term representation is given by:

$$\mathcal{R}(\mathcal{B}_c(x_u^{(n)}, b_u), \mathcal{B}_c(x_v^{(n)}, b_v)) = \begin{cases} \mathcal{R}_c(x_u^{(n)}, x_v^{(n)}) & \text{if } b_u = 0, b_v = 0 \\ \mathcal{R}_c(x_u^{(n)}, x_v^{(n)} + s) & \text{if } b_u = 0, b_v = 1 \\ \mathcal{R}_c(x_u^{(n)} + s, x_v^{(n)}) & \text{if } b_u = 1, b_v = 0 \\ \mathcal{R}_c(x_u^{(n)} + s, x_v^{(n)} + s) & \text{if } b_u = 1, b_v = 1 \end{cases} \quad (10)$$

(10) is submodular as $\mathcal{R}_c(x_u, x_v)$ is a L^{\natural} -convex function (since f is convex, its submodularity inequality $f(|x_u + s - x_v|) + f(|x_u - x_v - s|) \geq 2f(|x_u - x_v|)$ is always satisfied). The optimal convex move can be found with Murota's gradient descent algorithm [23]. It is worth noting that GC formulation does not impose any requirements on data fidelity term thus guaranteeing that the energy decreases. Hence, in this case the energy (1) is minimized but the optimal solution of multilabel problem is not secured.

4 Truncated convex prior algorithm

In this section, we present an effective method combining both moves introduced in Section 3 for minimizing energies with truncated convex prior functionals (2).

The convex move submodularity inequality is a function of (x_u, x_v) s.t. $u, v \in \mathcal{N}$ and $s \in \mathcal{S}$. The choice of \mathcal{S} influences the number of pairs of neighboring pixels $u \in \mathcal{V}$ which satisfies the convex move submodularity inequality. We examine the case where $\mathcal{S} = \{-1, +1\}$ and $f(x_u, x_v)$ is defined as in (2). To specify the sets of pixels the convex move applies to, we define \mathbb{T}_i for $0 \leq i \leq L$ to be the collection of all subsets $S_i^{\mathcal{V}}$ of \mathcal{V} such that $\forall \hat{u}, \hat{v} \in S_i^{\mathcal{V}}, |\hat{u} - \hat{v}| \leq i$. We note that all x_u belong to at least one $S_i^{\mathcal{V}}$ irrespective of i , and so the entire image is covered by \mathbb{T}_i . A convex move characterized by $\mathcal{S} = \{-1, +1\}$ is a function which maps \mathbb{T}_{T-1} onto \mathbb{T}_T , guaranteeing that the energy defined as (2) decreases with each move. This comes from the fact that the energy for the \mathbb{T}_{T-1} is represented exactly using our convex graph and as s is equal to either 1 or -1 , the solution belongs to \mathbb{T}_T .

Following [15], we define the edge capacities of graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$. The cost $c(u, v)$ between $(u, v) \in \mathcal{N}$ is set to $f(|x_u + s - x_v|) + f(|x_u - x_v - s|) - 2f(|x_u - x_v|)$ if $|x_u - x_v| < T$ and 0 otherwise. Because of the many such null connections, the final MRF is sparser which improves the time efficiency of the algorithm. The energy is guaranteed to go down, but the resulting labeling and corresponding energy are not as good as obtained by other minimizers. To improve our results, we combine this convex move with our proposed quantized move.

An arbitrary new labeling set by the quantized move part is not guaranteed to improve the energy with respect to the truncated convex prior energy (only a Potts-like energy is guaranteed to be minimized). However, we can easily impose this extra condition: the new labeling is accepted only if the proposed energy is better with respect to truncated convex prior energy, and rejected otherwise, which yields the desired effect. Since quantized move regularizes distant outliers,

it is a powerful complement method for convex moves, for which $\mathcal{S} = \{-1, +1\}$ regularizes close outliers. Now, we present our two-step algorithm alternating convex and quantized move. Here, $Q(x, k)$ denotes the quantized move of image x and interval k . We also denote the convex move by $C(x, s)$, where s is the considered step and x the input image. Note that the loops indexed by n and m are repeated until convergence.

Algorithm 1 (*Quantized-convex split moves algorithm*)

Fix $x^{(0)}, \mathcal{S} = \{-1, 1\}$

For $j = 0, 1, \dots$

$x^{(0)} = x^{(j)}$	For $n = 0, 1, \dots$												
	<table style="border-collapse: collapse; margin-left: 2em;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">Assign to $\bar{\mathcal{K}}$ a set of randomly ordered elements from \mathcal{K}</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">For $i = 0, 1, \dots, K$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;"> <table style="border-collapse: collapse; margin-left: 2em;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">Set k_i to be the i-th element of $\bar{\mathcal{K}}$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">$x = Q(x^{(n)}, k_i)$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">if $(E(x) \leq E(x^{(n)}))$ then $x^{(n+1)} = x$</td> </tr> </table> </td> </tr> </table>		Assign to $\bar{\mathcal{K}}$ a set of randomly ordered elements from \mathcal{K}		For $i = 0, 1, \dots, K$		<table style="border-collapse: collapse; margin-left: 2em;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">Set k_i to be the i-th element of $\bar{\mathcal{K}}$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">$x = Q(x^{(n)}, k_i)$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">if $(E(x) \leq E(x^{(n)}))$ then $x^{(n+1)} = x$</td> </tr> </table>		Set k_i to be the i -th element of $\bar{\mathcal{K}}$		$x = Q(x^{(n)}, k_i)$		if $(E(x) \leq E(x^{(n)}))$ then $x^{(n+1)} = x$
	Assign to $\bar{\mathcal{K}}$ a set of randomly ordered elements from \mathcal{K}												
	For $i = 0, 1, \dots, K$												
	<table style="border-collapse: collapse; margin-left: 2em;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">Set k_i to be the i-th element of $\bar{\mathcal{K}}$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">$x = Q(x^{(n)}, k_i)$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">if $(E(x) \leq E(x^{(n)}))$ then $x^{(n+1)} = x$</td> </tr> </table>		Set k_i to be the i -th element of $\bar{\mathcal{K}}$		$x = Q(x^{(n)}, k_i)$		if $(E(x) \leq E(x^{(n)}))$ then $x^{(n+1)} = x$						
	Set k_i to be the i -th element of $\bar{\mathcal{K}}$												
	$x = Q(x^{(n)}, k_i)$												
	if $(E(x) \leq E(x^{(n)}))$ then $x^{(n+1)} = x$												
$x^{(0)} = x^{(n)}$	For $m = 0, 1, \dots$												
	<table style="border-collapse: collapse; margin-left: 2em;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">Assign to $\bar{\mathcal{S}}$ a set of randomly ordered elements from \mathcal{S}</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">For $i = 0, 1$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;"> <table style="border-collapse: collapse; margin-left: 2em;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">Set s_i to be the i-th element of $\bar{\mathcal{S}}$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">$x^{(m+1)} = C(x^{(m)}, s_i)$</td> </tr> </table> </td> </tr> </table>		Assign to $\bar{\mathcal{S}}$ a set of randomly ordered elements from \mathcal{S}		For $i = 0, 1$		<table style="border-collapse: collapse; margin-left: 2em;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">Set s_i to be the i-th element of $\bar{\mathcal{S}}$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">$x^{(m+1)} = C(x^{(m)}, s_i)$</td> </tr> </table>		Set s_i to be the i -th element of $\bar{\mathcal{S}}$		$x^{(m+1)} = C(x^{(m)}, s_i)$		
	Assign to $\bar{\mathcal{S}}$ a set of randomly ordered elements from \mathcal{S}												
	For $i = 0, 1$												
	<table style="border-collapse: collapse; margin-left: 2em;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">Set s_i to be the i-th element of $\bar{\mathcal{S}}$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding-left: 5px;">$x^{(m+1)} = C(x^{(m)}, s_i)$</td> </tr> </table>		Set s_i to be the i -th element of $\bar{\mathcal{S}}$		$x^{(m+1)} = C(x^{(m)}, s_i)$								
	Set s_i to be the i -th element of $\bar{\mathcal{S}}$												
	$x^{(m+1)} = C(x^{(m)}, s_i)$												
$x^{(j+1)} = x^{(m)}$													

We now have our main result:

Proposition 2 *Algorithm 1 iteratively decreases the energy (1), with \mathcal{R} defined as a truncated convex function.*

Proof: This result comes straightforwardly from the previous discussion, where it was shown that all steps reduce the energy $E(x)$.

As this algorithm combines quantized and convex moves, it is important to understand what happens at the boundary between them. A difficulty is that neighboring pairs $u, v \in \mathcal{N}$ with labels $|x_u - x_v| = T$ cannot be represented exactly on the convex graph. This comes from the fact that the convex move cannot map \mathbb{T}_T to \mathbb{T}_{T-1} . We cope with this problem in a similar way as in [31], where α -expansions were shown to be able to minimize energies involving a truncated prior, as long as the number of pairs x_u, x_v not satisfying the submodularity inequality is relatively small. This is the reason why we limited the convex moves to $\mathcal{S} = \{-1, +1\}$. We represent truncated priors on convex move graph in a similar spirit.

5 Results

We implemented our proposed algorithm in the framework of the Middlebury MRF vision code (<http://vision.middlebury.edu/MRF/code/>), based on [21], so we could compare our approach with the following methods: ICM [13], α -expansion and swap moves [32, 6], MaxProdBP, BP-S (using software provided by Marshall Tappen [33]) and TRW-S [25, 24]. We also endeavoured to compare it with the range move, but range move did not work for our test because the value of T was too large. The tests were performed single-threaded on an Intel Xeon 2.5GHz with 32GB of RAM running RedHat Enterprise Linux 5.5. All algorithms were run either until full convergence for GC algorithms, ICM, and ours, or until the first oscillation for the other algorithms.

We evaluate our proposed algorithm only in the context of image restoration for different prior functions, namely truncated ℓ_2 and truncated ℓ_1 - ℓ_2 , defined as $\sqrt{\epsilon + x^2}$. In each case, we also examine the influence of parameter T . The grey scale images ($L = 255$) of size 512×512 (for ℓ_2) or 256×256 (for ℓ_1 - ℓ_2) were corrupted with additive zero mean Gaussian noise with standard deviation 25.3 corresponding to initial SNR values 13.75 dB, 15.09 dB, and 14.26 dB for images “gold rec”, “elaine”, and “barbara”, respectively. Consequently all experiments were performed with an ℓ_2 data fidelity term, which is most appropriate for this noise distribution. All the algorithms were initialized with an empty zero image. The algorithm accuracy is evaluated in terms of absolute error defined as $err = (E(x^*) - E(x_{TRW-S_i})) / E(x_{TRW-S_i})$, where $E(x_{TRW-S_i})$ is the lower bound value reported by TRW-S and $E(x^*)$ is an energy corresponding to the solution obtained by the algorithm. The restoration quality is evaluated in terms of SNR. The mean time, the energy, SNR, and the error presented in Table 1 and Table 2 are computed from 3 different realizations of the noise added to 3 considered images. The performance of our algorithm is also illustrated by energy vs. time plots (Fig. 2).

In Table 2, our *quantized-convex split moves* algorithm outperforms all other GC based algorithms in terms of minimum energy and time efficiency for truncated ℓ_2 prior. However, the best final energy is obtained by the BP (contrary to what was found in [21]) and TRW-S algorithms, the latter converging faster than the former. One can observe in Fig. 2 (a) that in the case of truncated ℓ_2 , TRW-S offers a speed/energy compromise comparable with our *quantized-convex split moves* algorithm when it is stopped early, for instance after two iterations. However, for truncated ℓ_1 - ℓ_2 prior, our algorithm is significantly faster (Fig. 2 (b)), while still achieving energies comparable with other algorithms (Table 2).

The quality of the results is also verified by inspecting the mean SNR value, which is not further improved by other algorithms in comparison to ours. Indeed, our algorithm appears to perform better at removing isolated noisy pixels compared with other algorithms (see Fig. 3(i)). Since our *Quantized-convex split moves* algorithm leads to very good results (Fig.3), is fast and less memory expensive than other algorithms, it appears to be well suited for image restoration application.

	$T = 25, \lambda = 2$			$T = 35, \lambda = 2$			$T = 50, \lambda = 1$		
	time	err	SNR	time	err	SNR	time	err	SNR
ICM	39.8	3.02×10^{-2}	20.11	39.06	1.71×10^{-2}	21.70	25.4	5.25×10^{-3}	21.53
BP-S	1807.8	8.29×10^{-4}	20.82	1658.7	5.10×10^{-4}	21.74	1641.8	7.21×10^{-5}	21.52
BP	153.0	1.05×10^{-3}	20.80	154.5	7.21×10^{-4}	21.76	153.3	1.17×10^{-4}	21.52
TRW-S	154.9	1.36×10^{-3}	20.85	154.6	7.55×10^{-4}	21.72	172.2	8.53×10^{-3}	21.54
α -exp	307.6	1.98×10^{-2}	20.53	294.6	2.01×10^{-2}	21.61	240.3	1.98×10^{-2}	21.28
$\alpha\beta$ swap	360.1	2.57×10^{-2}	20.33	362.3	1.48×10^{-2}	21.75	359.5	4.26×10^{-3}	21.53
Proposed	27.1	1.57×10^{-2}	21.53	28.6	6.30×10^{-3}	21.71	29.2	3.28×10^{-3}	21.51

Table 1. Truncated ℓ_2 prior results on 512×512 images. The SNR is given in dB, and the time in seconds. Best results are in bold. TRW-S and BP were stopped after 15 iterations (after this, the energy did not improve significantly).

	$T = 35, \lambda = 55$			$T = 50, \lambda = 45$			$T = 60, \lambda = 30$		
	time	err	SNR	time	err	SNR	time	err	SNR
ICM	104.4	2.69×10^{-2}	19.51	84.4	8.67×10^{-3}	20.57	46.8	2.39×10^{-3}	20.88
BP-S	4871.9	7.28×10^{-4}	20.08	5069.9	1.34×10^{-4}	20.68	3866.7	1.42×10^{-5}	20.97
BP	13950.0	9.40×10^{-4}	20.12	16048.7	2.10×10^{-4}	20.69	14902.3	4.31×10^{-5}	20.97
TRW-S	2508.9	2.54×10^{-4}	20.06	2259.0.4	4.30×10^{-5}	20.66	2852.2	5.08×10^{-6}	20.97
α -exp	61.8	7.96×10^{-3}	20.10	50.3	7.36×10^{-3}	20.72	50.9	7.37×10^{-3}	20.80
$\alpha\beta$ swap	200.4	1.12×10^{-2}	19.94	178.9	4.31×10^{-3}	20.63	112.7	1.19×10^{-3}	20.95
Proposed	9.4	1.16×10^{-2}	20.37	9.3	4.00×10^{-3}	20.86	7.9	1.51×10^{-3}	21.19

Table 2. Truncated ℓ_1 - ℓ_2 prior results with $\epsilon = 10$ on 256×256 images. The SNR is given in dB, and the time in seconds. Best results are in bold.

6 Conclusion and future work

In this paper, we have presented a novel move-based algorithm to solve GC problems with truncated convex priors in the context of image denoising. Our move is split in two parts, a first Potts-like move that denoises a quantized version of the image, and a second move that processes the result of the first move according to a fully convex prior. We have shown that combining these moves corresponds to denoising with a truncated convex prior. For a convex prior truncated at threshold T and for an image with L labels, the Potts-like denoising operates on L/T labels and the convex part on T labels only. This results in two optimizations over a much reduced set of labels for most useful values of T , and therefore it translates into large savings in computing time. Because only submodular moves are effected, the algorithm is guaranteed to converge in finite time. The result of these moves appears better in terms of energy than all moves, and depending on the problems, our algorithm is at least 5 times and up to several orders of magnitude faster than current state-of-the-art algorithms. We believe this constitutes an interesting compromise between efficiency and precision.

Since we use barely modified versions of Potts optimization and discrete convex optimization methods, future progress in this area will also translate into improvements for the proposed method. In particular, future work will include analyzing primal-dual methods for convex optimization. Precision can also be improved by using more sophisticated Potts-like moves. We will also explore




Fig. 2. Energy versus log time characteristics of convergence for algorithm comparison: ICM (yellow), BP-S (red), BP (green), TRW-S (cyan), $\alpha\beta$ swap (magenta), α -exp (blue), ours (black). (a,b) illustrates the case of ℓ_2 and ℓ_1 - ℓ_2 prior, respectively.

the behaviour of our algorithm with non-convex data terms and consider other applications, such as stereo-vision.

References

1. Veksler, O.: Efficient graph-based energy minimization methods in computer vision. PhD thesis, Cornell University, Ithaca, NY, USA (1999)
2. Huang, J., Mumford, D.: Statistics of natural images and models. In: Computer Vision and Pattern Recognition, 1999. IEEE Computer Society Conference on Computer Vision, (Fort Collins, CO, USA)
3. Nikolova, M.: Minimizers of cost-functions involving non-smooth data-fidelity terms. application to the processing of outliers. *SIAM J. on Numerical Analysis* **40** (2002) 965–994
4. Boykov, Y., Veksler, O., Zabih, R.: Markov random fields with efficient approximations. In: CVPR. (1998) 648–655
5. Ishikawa, H.: Exact optimization for Markov random fields with convex priors. *IEEE Transaction on Pattern Analysis and Machine Intelligence* **25** (2003) 1333–1336
6. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. *IEEE Transaction on Pattern Analysis and Machine Intelligence* **23** (2001) 1222–1239
7. Woodford, O.J., Torr, P.H.S., Reid, I.D., Fitzgibbon, A.W.: Global stereo reconstruction under second order smoothness priors. In: IEEE Conference on Computer Vision and Pattern Recognition. (2008)
8. Sinha, S.N., Mordohai, P., Pollefeys, M.: Multi-view stereo via graph cuts on the dual of an adaptive tetrahedral mesh. In: Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on, (Rio de Janeiro) 1–8
9. Xiao, J., Shah, M.: Motion layer extraction in the presence of occlusion using graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27** (2007) 1644–1659
10. Boykov, Y.Y., Jolly, M.P.: Interactive graph cuts for optimal boundary & region segmentation of objects in n-d images. In: Computer Vision, 2001. ICCV 2001.




Fig. 3. The image restoration results for truncated ℓ_1 - ℓ_2 prior with threshold $T = 50$, ℓ_2 data fidelity term, and $\lambda = 2$.

Proceedings. Eighth IEEE International Conference on, (Vancouver, BC , Canada) 105–112

11. Darbon, J., Sigelle, M.: Image restoration with discrete constrained total variation part ii: Levelable functions, convex priors and non-convex cases. *JMIV* **26** (2006) 277 – 291
12. Geman, S., Geman, D.: Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *tpami* **6** (1984) 721–741
13. Besag, J.: On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)* **48** (1986) 259–302
14. Murota, K.: *Algorithms in discrete convex analysis.* (2000)
15. Kolmogorov, V., Zabih, R.: What energy functions can be minimized via graph cuts? *IEEE Transaction on Pattern Analysis and Machine Intelligence* **26** (2004) 147–159

16. Schlesinger, D. and Flach, B.: Transforming an arbitrary minsum problem into a binary one. Technical report, Dresden University of Technology (2008)
17. Veksler, O.: Graph cut based optimization for MRFs with truncated convex priors. In: IEEE Conference on Computer Vision and Pattern Recognition, Minneapolis, MN (2007) 1–8
18. Kumar, M.P., Torr, P.H.S.: Improved moves for truncated convex models. In: Proceedings of Advances in Neural Information Processing Systems. (2008)
19. Lempitsky, V., Rother, C., Roth, S., Blake, A.: Fusion moves for markov random field optimization. IEEE Transactions on Pattern Analysis and Machine Intelligence **32** (2010) 1392–1405
20. Veksler, O.: Multi-label moves for MRFs with truncated convex priors. In: The 7th International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition EMMCVPR'09. (2009) 1–8
21. Szeliski, R., Zabih, R., Scharstein, D., Veksler, O., Kolmogorov, V., Agarwala, A., Tappen, M., Rother, C.: A comparative study of energy minimization methods for markov random fields with smoothness-based priors. IEEE Transactions on Pattern Analysis and Machine Intelligence **30** (2008) 1068–1080
22. Kumar, M.P., Koller, D.: MAP estimation of semi-metric MRFs via hierarchical graph cuts. In: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence. UAI '09, Arlington, Virginia, United States, AUAI Press (2009) 313–320
23. Murota, K.: On steepest descent algorithms for discrete convex functions. SIAM Journal on Optimization **14** (2004) 699–707
24. Kolmogorov, V.: Convergent tree-reweighted message passing for energy minimization. IEEE Transactions on Pattern Analysis and Machine Intelligence **28** (2006) 1568–1583
25. Wainwright, M.J., Jaakkola, T.S., Willsky, A.S.: MAP estimation via agreement on trees: message-passing and linear programming. IEEE Transactions on Information Theory **51** (2005) 3697–3717
26. Kolmogorov, V., Rother, C.: Comparison of energy minimization algorithms for highly connected graphs. In: proceedings of ECCV. Volume II. (2006) 1–15
27. Felzenszwalb, P.F., Huttenlocher, D.R.: Efficient belief propagation for early vision. (In: Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision) 261–268
28. Gupta, A., Tardos, .: A constant factor approximation algorithm for a class of classification problems. In: In Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science. (2000) 333–342
29. Ford, J.L.R., Fulkerson, D.R.: Flows in Networks. Princeton University Press, Princeton, NJ (1962)
30. Kolmogorov, V., Shioura, A.: New algorithms for convex cost tension problem with application to computer vision. Discrete Optimization **6** (2009) 378–393
31. Rother, C., Kumar, S., Kolmogorov, V., Blake, A.: Digital tapestry [automatic image synthesis]. In: Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on. (2005) 589–596
32. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. IEEE Transaction on Pattern Analysis and Machine Intelligence **26** (2004) 1124–1137
33. Tappen, M.F., Freeman, W.T.: Comparison of graph cuts with belief propagation for stereo, using identical MRF parameters. In: Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on, Nice, France (2008) 900–906