

A Convex Discrete-Continuous Approach for Markov Random Fields

Christopher Zach and Pushmeet Kohli

Microsoft Research Cambridge, UK

Abstract. We propose an extension of the well-known LP relaxation for Markov random fields to explicitly allow continuous label spaces. Unlike conventional continuous formulations of labelling problems which assume that the unary and pairwise potentials are convex, our formulation allows them to be general piecewise convex functions with continuous domains. Furthermore, we present the extension of the widely used efficient scheme for handling L^1 smoothness priors over discrete ordered label sets to continuous label spaces. We provide a theoretical analysis of the proposed model, and empirically demonstrate that labelling problems with huge or continuous label spaces can benefit from our discrete-continuous representation.

1 Introduction

Energy minimization algorithms have been used to infer the Maximum a Posteriori (MAP) solutions of probabilistic formulations of many labeling problems encountered in Computer Vision. For instance, the Graph Cut algorithm, which enables the exact minimization of certain classes of pseudoboolean functions, has been used for building extremely efficient interactive tools for image segmentation [1]. Researchers, however, have not been able to replicate the success of discrete energy minimization algorithms, such as Graph Cuts and Message Passing, for building fast methods for solving labeling problems with very large or even continuous label sets such as image denoising and optical flow.

The above-mentioned problems, like many others encountered in Computer Vision, are defined over variables taking continuous labels. If the data term and the smoothness costs are convex functions, the underlying Markov random field (MRF) can be solved directly with convex optimization methods. However, this is not the case for many real world image labelling problems which have non-convex data terms and sometimes non-convex smoothness priors. For handling such problems, standard MRF algorithms discretize the label space, sample the energy costs associated with the chosen discrete label set, and then minimize the resulting energy which is defined over discrete variables.

The label discretization needed for handling continuous label problems can lead to arbitrarily large deviations from the true solution. The number of discrete labels needed to represent a continuous label set should to be large enough to provide a faithful solution to the original, continuous problem. Increasing the

number of labels implies substantial memory requirements, e.g. the memory consumption grows at least linearly with the number of labels, but can also grow quadratically (depending on the utilized formulation and optimization method). In particular, this often rules out e.g. leveraging graphics processing units to accelerate MAP estimation in such MRFs.

In this paper, we propose a convex discrete-continuous formulation of labeling problems. Our approach can compactly represent arbitrary data costs while still operating on a continuous label space. In this sense, it brings together the best of both discrete and continuous formulations. It is inspired from the observation that the data cost can be approximated as a lower envelope of convex basis functions. Fig. 1 illustrates the difference between (unary) potentials handled by well-established LP relaxations for MRFs and the potentials allowed in our approach. In this work the focus lies on how a more expressive potential function Fig. 1(b) can be used to efficiently approximate an arbitrary cost profile for a multi-label problem.




Fig. 1. Unary potentials as induced by standard LP relaxations (left) under a *special ordered set of type 2* interpretation [2], and one example that can be modeled using our discrete-continuous approach (right)

An outline of the paper follows. We discuss prior work on inference in MRFs with a focus on continuous label sets in section 2. Section 3 provides the mathematical notations and presents the main result enabling our convex discrete-continuous formulation for labeling problems, which is described in detail in section 4. We compare the performance of our method with traditional discrete optimization methods in section 5, and conclude in section 6 by listing some key conclusions and ideas for future work.

2 Related Work

There is a vast literature on different formulations and solution methods for Markov random fields. We refer to [3] for an extensive treatment of the respective statistical background, and e.g. [4] for an exposition of the corresponding LP relaxations. In certain cases the LP relaxation of discrete labeling problems provides globally optimal answers as discussed in [5,6] for L^1 smoothness priors and in [7] for general pairwise potentials that are convex in the label difference.

The key observation that motivates our approach is, that unary and pairwise clique potentials for many continuous labeling problems encountered in the real world are not arbitrarily complex, but can be represented or at least well-approximated by piecewise convex functions¹. Our proposed formulation generalizes the one described in [8], which completely convexifies the unary potentials in order to obtain (in combination with a convex smoothness prior) a convex problem over continuous labels. This method is a consequent continuation of the celebrated local linearization/convexification principle prominently utilized in the computation of optical flow (dating back to [9,10]). The assumption of parametric or otherwise structured potentials in general also enables efficient inference methods as discussed in [11,12,13], where the focus lies on exploiting structured representation of higher-order clique potentials.

Continuous label spaces naturally raise the question of also using a continuous representation of the underlying domain (e.g. the 2D image domain). Recently, the analysis of the Mumford-Shah functional and its extensions in [14] proved to be a source of inspiration to address general labeling problems on image grids (e.g. [15]). [16,17] propose to solve discretized versions of the underlying continuous model [14], thereby also handling continuous transitions of labels between neighboring pixels. Their formulation is based on a primal-dual saddlepoint problem and the underlying representation turns out to be substantially different to our proposed one.²

Peng et al [18] share their goal with the one of our work—inference in general MRFs with continuous label spaces—, but employ a completely different representation and inference procedure, which is based on convex belief propagation and particle sampling to represent beliefs over continuous labels. Finally, we want to mention that there are algorithmically very different approaches to tackle large or continuous label spaces, e.g. by adaptive refinement of the discrete label set as proposed in [19].

3 Notations and Preliminary Result

3.1 Notations

In the following we consider extended real-valued functions $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$. The domain of f , $\text{dom}(f)$, is $\{x \in \mathbb{R}^n : f(x) < \infty\}$, and we assume that $\text{dom}(f) \neq \emptyset$. Hence, constraints on the feasible domain and infinite function values can be interchanged. The convex conjugate of f , denoted by f^* , is defined as $f^*(y) = \sup_x x^T y - f(x)$. The biconjugate f^{**} is obtained by applying convex conjugation twice. It is known that for any function f the convex conjugate f^* is a lower-semicontinuous (l.s.c.) convex function, and $f^{**} = f$ iff f is convex and l.s.c. Otherwise f^{**} is the lower convex envelope of f , i.e. the supremum of all convex functions below the epigraph of f .

¹ This is also the main assumption of our framework.

² An in-depth discussion of the relation between the models is subject of an extended manuscript.

For a convex set C we will use ι_C for its indicator function. i.e. $\iota_C(x) = 0$ for $x \in C$ and ∞ otherwise. We will also write e.g. $\iota\{f(x) \leq 0\}$ to convert a constraint (here $f(x) \leq 0$) into a function. For a convex function f the corresponding perspective will be denoted by f_\circ . It is defined indirectly as the convex conjugate $f_\circ = (f_\circ)^{**}$ of

$$(f_\circ)^*(z, w) = \iota\{-z \geq f^*(w)\}.$$

For $x > 0$ we have $f_\circ(x, y) = xf(y/x)$. We use the somewhat indirect definition above to handle also the case $x = 0$. Since $(f_\circ)^*$ is an indicator function (i.e. a set of constraints), f_\circ is convex and positively 1-homogeneous, i.e. $f_\circ(kx, ky) = kf_\circ(x, y)$. Again, for $x > 0$ this can be verified directly from $kxf(ky/kx) = kxf(y/x)$, but this extends also to $x = 0$. If the domain of f^* is \mathbb{R} (or \mathbb{R}^D if f takes a vector argument), one obtains $f_\circ(0, y) = \iota\{y = 0\}$ (we refer to the supplementary material).

In this work we focus on MRFs with pairwise interactions. Hence, the potentials will be defined on nodes $s \in \mathcal{V}$ and edges $(s, t) \in \mathcal{E}$ of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. We will use $\sum_{s \sim t}$ as the easier-to-read replacement for $\sum_{s,t:(s,t) \in \mathcal{E}}$.

3.2 Minimizing a Family of Convex Functions

It is well known and has been extensively used in the literature, that the minimum of a finite number of values, e.g. $\{\theta^1, \dots, \theta^n\}$ can be formulated as convex problem

$$\min_i \theta^i = \min_{\mathbf{x}} \sum_{i=1}^n \theta^i x^i = \min_{\mathbf{x}} \theta^T \mathbf{x} \tag{1}$$

subject to $x_i \geq 0, \sum_{i=1}^n x_i = 1$ (i.e. $\mathbf{x} \in \Delta_n$, the n -dimensional unit simplex). Note that the right hand side is a convex problem in \mathbf{x} . The equivalence above is essentially the basis of LP relaxations of Markov random fields, where x is usually called the pseudo-marginal. Introducing the explicit weight vector \mathbf{x} allows to add constraints linking unary and clique potentials. We extend this construction to a family of convex functions. Let $\{f^1, \dots, f^n\}$ be a family of convex functions, $f^i : [0, 1] \rightarrow \mathbb{R}$ and we consider the minimization problem

$$\min_i \min_{\xi \in [0,1]} f^i(\xi). \tag{2}$$

The task is to reformulate this problem into a form explicitly incorporating a convex combination of weights, similar to Eq. 1. We will use the following reformulation in our construction:

Observation 1. *We have that*

$$\min_i \min_{\xi \in [0,1]} f^i(\xi) = \min_{\mathbf{x} \in \Delta^n, \mathbf{y}: y_i \leq x_i} \sum_{i=1}^n f_\circ^i(x_i, y_i) = \min_{\mathbf{x} \in \Delta^n, \mathbf{y}: y_i \leq x_i} \sum_{i=1}^n x_i f^i(y_i/x_i). \tag{3}$$

The proof is given in the supplementary material. At this point we want to add a few remarks:

- If $f^i(\xi) = \theta^i$ (constant functions), then Eq. 3 degenerates to Eq. 1. Surprisingly, the same is true if f^i linearly interpolates between e.g. θ^i and θ^{i+1} , i.e. $f^i(\xi) = \xi\theta^{i+1} + (1 - \xi)\theta^i$ (like in Fig. 1(a)). This follows from the fact, that the minimum is (also) attained at one breakpoint for piecewise linear functions.
- If $p := \min_i \min_{\xi \in [0,1]} f^i(\xi)$ is not unique (with respect to i), then the minimizer \mathbf{x}^* is not strictly binary, and for all i with $x_i^* > 0$, one has $x_i^* f_i(y_i^*/x_i^*) = p$ for optimal $(\mathbf{x}^*, \mathbf{y}^*)$.

This construction also extends to families of functions $\{f^i\}_{i=1,\dots,n}$ with vector-valued arguments, i.e. $f^i : [0, 1]^D \rightarrow \mathbb{R}$ with D being the dimension of the domain. In this setting we have an analogy to Eq. 3,

$$\min_i \min_{\xi \in [0,1]^D} f^i(\xi) = \min_{\mathbf{x} \in \Delta^n, \mathbf{Y}: \mathbf{y}_i \leq x_i} \sum_{i=1}^n f_{\mathcal{O}}^i(x_i, \mathbf{y}_i), \tag{4}$$

where $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n) \in [0, 1]^{D \times n}$ is now a $D \times n$ matrix instead of a n -vector, and $\mathbf{y}_i \leq x_i$ is understood element-wise.

4 Convex Discrete-Continuous Formulation for MRFs

This section describes our combined discrete-continuous approach to represent labels for inference in MRFs. First, we review and reinterpret the standard local polytope relaxation for MRFs, and subsequently introduce our extensions.

Recall from section 1 that we are addressing continuous labeling problems, where a label value from the range $[0, L]$ is assigned to each node s , where L is the number of segments to represent the piecewise convex potentials. We are going to represent the continuous label l_s assigned at node s as $l_s = \sum_i (ix_s^i + y_s^i)$, where $x_s^i \in \{0, 1\}$ indicates the selected label range $[i, i + 1]$ and $y_s^i \in [0, 1]$ is the continuous offset within the subrange. Since exactly one subrange needs to be selected, we have $\sum_i x_s^i = 1$. Enforcing x_s^i to be either 0 or 1 is hard, hence we relax $x_s^i \in \{0, 1\}$ to $x_s^i \in [0, 1]$. The restriction to label ranges $[0, L]$ is not limiting, since general bounded label domains can be scaled and translated without essentially changing the minimizer.

4.1 MAP Inference Using LP-MRF Relaxations Revisited

If one recalls the standard LP relaxation of discrete and labeling problem with at most pairwise potentials, assigning one out of L states to each node,

$$E_{\text{LP-MRF}}(\mathbf{x}) = \sum_{s,i} \theta_s^i x_s^i + \sum_{s \sim t} \sum_{i,j} \theta_{st}^{ij} x_{st}^{ij} \tag{5}$$

subject to $x_s \in \Delta^L$, $x_{st} \in \Delta^{L^2}$, and the marginalization constraints,

$$\sum_j x_{st}^{ij} = x_s^i, \quad \sum_i x_{st}^{ij} = x_t^j \quad \forall s \sim t, \forall i, j, \tag{6}$$

then it is immediately clear that the objective is a sum of convex combinations of the clique potentials θ_s^i and θ_{st}^{ij} , respectively. In the general higher-order MRF setting there is one such convex combination for each clique. Since the ‘‘barycentric’’ weights (or pseudo-marginals), x_s^i and x_{st}^{ij} , are made explicit in the above formulation, they can be coupled via marginalization constraints, thereby connecting clique variables x_{st}^{ij} with node marginals x_s^i . By recalling the extension described in Section 3.2 to find the global minimum of finite families of convex functions via explicit weights, it becomes imminent that E_{LP-MRF} can be generalized in order to allow piecewise convex potentials.

4.2 Unary Potentials

The cost induced by unary potentials, $\sum_i \theta_s^i x_s^i$, is easy to generalize: we assume that real-valued convex functions $f_s^i : [0, 1] \rightarrow \mathbb{R}$ are given, then the new unary cost reads as

$$\sum_s \sum_i (f_s^i)_{\circlearrowleft} (x_s^i, y_s^i) = \sum_s \sum_i x_s^i f_s^i (y_s^i / x_s^i),$$

where we have also introduced continuous variables $y_s^i \geq 0$ to represent the minimizer of f_s^i . Since $\text{dom}(f_s^i) = [0, 1]$ (i.e. $f_s^i(\xi) = \infty$ for $\xi \notin [0, 1]$), we have additional implicit constraints $y_s^i \leq x_s^i$ (in addition to the unit simplex constraint on x_s). Since an assignment of an integral label value $i > 0$ at node s is represented twice (as $x_s^i = 1, y_s^i = 0$, and alternatively as $x_s^{i-1} = 1, y_s^i = 1$), we require that $f_s^i(0) = f_s^{i-1}(1)$ (recall the graph in Fig. 1(b), which is continuous).

4.3 The General Discrete-Continuous Model

Generalizing the pairwise cost to continuous labels is far less trivial. In the following we present two models: the first formulation described in this section is the proper generalization of the LP relaxation Eq. 5 and introduces extended marginalization constraints. The second model discussed in Section 4.4 specifically addresses L^1 smoothness costs, and is much more efficient to optimize.

Assume that a family of convex functions $f_{st}^{ij} : [0, 1]^2 \rightarrow \mathbb{R}$ is given. Like for the unary potentials we require that the induced graph of the smoothness cost is continuous, i.e. $f_{st}^{ij}(0, \cdot) = f_{st}^{i-1, j}(1, \cdot)$ and $f_{st}^{ij}(\cdot, 0) = f_{st}^{i, j-1}(\cdot, 1)$. We replace the standard pairwise cost for edge (s, t) ,

$$\sum_{i, j} \theta_{st}^{ij} x_{st}^{ij}, \quad \text{by} \quad \sum_{i, j} (f_{st}^{ij})_{\circlearrowleft} (x_{st}^{ij}, y_{st}^{\rightleftharpoons ij}) = \sum_{i, j} x_{st}^{ij} f_{st}^{ij} \left(\frac{y_{st \rightarrow s}^{ij}}{x_{st}^{ij}}, \frac{y_{st \rightarrow t}^{ij}}{x_{st}^{ij}} \right),$$

where we introduced local copies $y_{st \rightarrow s}^{ij}$ and $y_{st \rightarrow t}^{ij}$ of y_s^i and y_t^j only relevant in the scope of a label transition from i to j . Further, we use $y_{st \rightleftharpoons}^{ij}$ for the 2-vector $(y_{st \rightarrow s}^{ij}, y_{st \rightarrow t}^{ij})$. The values of $y_{st \rightarrow s}^{ij}$ and $y_{st \rightarrow t}^{ij}$, respectively, must be consistent with the node values, y_s^i and y_t^j : we propose the following “decomposition” constraints,

$$y_s^i = \sum_j y_{st \rightarrow s}^{ij} \quad \text{and} \quad y_t^j = \sum_i y_{st \rightarrow t}^{ij}. \quad (7)$$

These constraints are not arbitrary, but define (together with the marginalization constraints on x_{st} and the bounds constraints $0 \leq y_{st \rightarrow s}^{ij} \leq x_{st}^{ij}$, $0 \leq y_{st \rightarrow t}^{ij} \leq x_{st}^{ij}$) the convex hull of allowed configurations. We refer to the supplementary material for the derivation. By combining the unary costs from Section 4.2 and the extension for pairwise costs described above, one obtains the following discrete-continuous labeling model,

$$E_{\text{DC-MRF}}(\mathbf{x}, \mathbf{y}) = \sum_{s,i} (f_s^i)_{\circlearrowleft} (x_s^i, y_s^i) + \sum_{s \sim t} \sum_{i,j} (f_{st}^{ij})_{\circlearrowleft} (x_{st}^{ij}, y_{st \rightleftharpoons}^{ij}) \quad (8)$$

$$\text{s.t. } \sum_j x_{st}^{ij} = x_s^i, \sum_i x_{st}^{ij} = x_t^j, y_s^i = \sum_j y_{st \rightarrow s}^{ij}, y_t^j = \sum_i y_{st \rightarrow t}^{ij}$$

$$x_s \in \Delta^L, x_{st} \in \Delta^{L^2}, 0 \leq y_s^i \leq x_s^i, 0 \leq y_{st \rightleftharpoons}^{ij} \leq x_{st}^{ij}.$$

The constraints $0 \leq y_s^i \leq x_s^i$ and $0 \leq y_{st \rightleftharpoons}^{ij} \leq x_{st}^{ij}$ are actually redundant, since these constraints are implied by restricting f_s^i and f_{st}^{ij} to have domain $[0, 1]$ and $[0, 1]^2$, respectively. Since the underlying implicit domain constraints might be easily missed, we state these constraints on \mathbf{y} here explicitly.

By comparing $E_{\text{LP-MRF}}$ (Eq. 5) with $E_{\text{DC-MRF}}$ it is clear that the latter one is a proper generalization of the LP relaxation and reduces to $E_{\text{LP-MRF}}$ for piecewise linear potentials f_s^i and f_{st}^{ij} . The downside of this model is, that it is expensive to optimize. Unlike $E_{\text{LP-MRF}}$, which can be efficiently optimized by dual coordinate methods (usually referred as message passing [20,4,21,22]), the dual of $E_{\text{DC-MRF}}$ is more involved and may have $O(L^2)$ variables per node (depending whether the $0 \leq y_{st \rightleftharpoons}^{ij} \leq x_{st}^{ij}$ constraints are moved into f_{st}^{ij} or not).

It is natural to ask whether the model in Eq. 8 returns globally optimal results for arbitrary unary potentials and convex pairwise priors as it is the case for the LP relaxation Eq. 5 [7]. Unfortunately, the answer seems to be negative as indicated by our experiments. Global optimality is currently guaranteed only for piecewise linear unary potentials and piecewise linear and convex pairwise priors, and to globally convex unary and pairwise priors, i.e. the original problem in terms of continuous labels assigned to nodes is convex. While in such cases it is not reasonable to artificially subdivide the label range into $L > 1$ segments and to optimize $E_{\text{DC-MRF}}$, it is at least encouraging that the minimizer of $E_{\text{DC-MRF}}$ returns the same solution as the original convex problem. Hence, the model in Eq. 8 does not weaken the optimal solution.

We will give a sketch of the proof: let $E_1(\mathbf{z}) = \sum_s F_s(\mathbf{z}) + \sum_{s \sim t} G_{st}(\mathbf{z})$ be the original problem, with F_s and G_{st} convex functions corresponding to the

unary and pairwise costs, respectively. For a given \mathbf{z} we can determine feasible arguments for $E_{\text{DC-MRF}}$ via e.g.

$$x_s^i(\mathbf{z}) = \begin{cases} 1 & \text{if } z_s \in [i, i + 1) \\ 0 & \text{otherwise} \end{cases} \quad y_s^i(\mathbf{z}) = \begin{cases} z_s - i & \text{if } z_s \in [i, i + 1) \\ 0 & \text{otherwise.} \end{cases}$$

Further, $f_s^i(\xi) = F_s(\xi + i)$ and $f_{st}^{ij}(\xi_s, \xi_t) = G_{st}(\xi_s + i, \xi_t + j)$. Let \mathbf{z}^* a minimizer of E_1 and $(\mathbf{x}^*, \mathbf{y}^*)$ the minimizer of the corresponding discrete-continuous program $E_{\text{DC-MRF}}$ Eq. 8. Since \mathbf{z}^* can be converted into a feasible argument $(\mathbf{x}(\mathbf{z}^*), \mathbf{y}(\mathbf{z}^*))$ of $E_{\text{DC-MRF}}$, we have $E_1(\mathbf{z}^*) \geq E_{\text{DC-MRF}}(\mathbf{x}^*, \mathbf{y}^*)$. On the other hand, $E_{\text{DC-MRF}}$ consists only of convex combinations (subject to marginalization constraints) of samples from convex functions, hence by Jensen’s inequality we have (for the unary potentials)

$$\sum_i (f_s^i)_{\circlearrowleft}(x_s^i, y_s^i) = \sum_i x_s^i f_s^i(y_s^i/x_s^i) \geq F_s \left(\sum_i (ix_s^i + y_s^i) \right).$$

A similar inequality holds for the pairwise terms. Overall we obtain $E_1(\mathbf{z}^*) \leq E_{\text{DC-MRF}}(\mathbf{x}^*, \mathbf{y}^*)$ and together with $E_1(\mathbf{z}^*) \geq E_{\text{DC-MRF}}(\mathbf{x}^*, \mathbf{y}^*)$ equality holds. Experimentally we confirm this fact in Fig. 2, which shows the results of image denoising using the (anisotropic) Rudin-Osher-Fatemi (ROF) model [23],

$$E_{\text{ROF}}(u; f) = \frac{\lambda}{2} \sum_s (u_s - f_s)^2 + \sum_{s \sim t} |u_s - u_t|. \tag{9}$$

Since the data term in the ROF energy is quadratic, f_s^i are quadratic and given by $f_s^i(\xi) = \lambda(i + \xi - f_s)^2/2$ with $\xi \in [0, 1]$. The L^1 smoothness term is slightly more involved:

$$f_{st}^{ij}(\xi_s, \xi_t) = \begin{cases} |\xi_s - \xi_t| & \text{if } i = j \\ \pm(i - j + \xi_s - \xi_t) & \text{if } i \neq j \end{cases} \tag{10}$$

If e.g. $i > j$ we have $z_s = i + \xi_s > z_t = j + \xi_t$, and the L^1 smoothness cost is therefore $|z_s - z_t| = i - j + \xi_s - \xi_t$. The other cases are similar.

4.4 Generalizing the “battleship” construction for L^1 Smoothness

The LP relaxation of standard discrete labeling problems with general pairwise smoothness costs (recall Eq. 5) requires to maintain $O(L^2)$ variables per node (with respect to the label subrange count L) in the primal domain (the pairwise pseudo-marginals x_{st}^{ij}). The constructions presented in [5,6]³ maintain only $O(L)$ unknowns per node. In order to understand the construction for L^1 smoothness in our discrete-continuous approach, the properties of Eq. 5 must be analyzed for this particular choice of pairwise potentials. The corresponding LP relaxation reads as

³ Which is sometimes called the “battleship” construction due to the graph structure.



(a) $E_{ROF} = 143992$ (b) $E_{ROF} = 144002$ (c) $E_{ROF} = 565809$ (d) $E_{ROF} = 565973$

Fig. 2. Visual and energy comparison between direct optimization [24] of the ROF image denoising model (a, c), and artificially subdividing the label space $[0, 255]$ into 4 even ranges and subsequent optimization of E_{DC-MRF} (b, d). The small differences in the obtained energy are likely due to the ultimately slow convergence of minimizing E_{DC-MRF} and due to numerical inaccuracies in require cube root computations.

$$E_{LP-MRF-L^1}(\mathbf{x}) = \sum_{s,i} \theta_s^i x_s^i + \sum_{s \sim t} \sum_{i,j} |i - j| x_{st}^{ij} \tag{11}$$

subject to $x_s \in \Delta^L$, $x_{st} \in \Delta^{L^2}$, and the marginalization constraints Eq. 6. The constraint $\sum_{i,j} x_{st}^{ij} = 1$ is redundant and can be dropped, thus $x_{st} \in \Delta^{L^2}$ can be replaced by $x_{st}^{ij} \geq 0$. The link between Eq. 11 and the battleship construction is, that x_{st}^{ii} is eliminated from the marginalization constraints and the corresponding non-negativity constraint is dropped.⁴

Via Fenchel duality it can be shown that the capacity constraints in the dual of Eq. 11 are given by $p_{st}^j + p_{ts}^j \leq |i - j|$ and $p_{st}^i - p_{ts}^i = 0$, respectively, for dual variables p_{st}^i, p_{ts}^i , representing the directed flow over edge (s, t) . Since $p_{st}^i - p_{ts}^i = 0$, p_{ts}^i can be eliminated, leaving capacity constraints $p_{st}^i - p_{st}^j \leq |i - j|$. Since $|\cdot|$ is a metric, all other capacity constraints follow from $p_{st}^i - p_{st}^{i+1} \leq 1$ and $p_{st}^i - p_{st}^{i-1} \leq 1$, i.e. only these constraints are necessary. By going back to the primal program, this translates to the following objective,

$$E_{LP-Linear}(\mathbf{x}) = \sum_{s,i} \theta_s^i x_s^i + \sum_{s \sim t} \sum_i \left(x_{st}^{i,i+1} + x_{st}^{i,i-1} \right) \tag{12}$$

$$\text{s.t. } x_s^i = x_{st}^{ii} + x_{st}^{i,i+1} + x_{st}^{i,i-1} \quad x_t^i = x_{st}^{ii} + x_{st}^{i-1,i} + x_{st}^{i+1,i}$$

$$x_s \in \Delta^L, x_{st}^{i,i+1} \geq 0, x_{st}^{i,i-1} \geq 0.$$

It is possible to further simplify Eq. 12 to obtain the minimum-cut formulation presented in [5,6], but the above formulation is the basis for our construction.

Before we continue, it is instructive to see, how a jump e.g. from label i to $i+3$ at edge (s, t) is actually represented in the above program: one has $x_{st}^{i,i+1} = 1$, $x_{st}^{i+1,i+1} = -1$, $x_{st}^{i+1,i+2} = 1, x_{st}^{i+2,i+2} = -1$ and finally $x_{st}^{i+2,i+3} = 1$, leading

⁴ Dropping $x_{st}^{ii} \geq 0$ converts general pairwise potentials θ_{st}^{ij} into a metric, hence nothing is lost whenever θ_{st}^{ij} is already a metric. A more detailed discussion of these relations can be found in [25].

the the correct smoothness cost of 3 and still satisfying all marginalization constraints. Thus, label discontinuities along an edge (s, t) correspond to chains of $+1$ for $x_{st}^{i,i+1}$ (or $x_{st}^{i,i-1}$) and -1 's at respective x_{st}^{ii} to fulfill the marginalization constraints for in-between labels.

In view of the energy Eq. 12 we only need to specify pairwise potentials for transitions x_{st}^{ii} , $x_{st}^{i,i+1}$, and $x_{st}^{i,i-1}$. These potentials are similar to the ones given in Eq. 10 and read as

$$f_{st}^{ii}(\xi_s, \xi_t) = |\xi_s - \xi_t| \qquad f_{st}^{i,i\pm 1}(\xi_s, \xi_t) = 1 \pm (\xi_t - \xi_s). \tag{13}$$

Further, since x_{st}^{ii} can be negative, we split x_{st}^{ii} into non-negative and non-positive components, i.e. we replace x_{st}^{ii} by $x_{st}^{ii} - \bar{x}_{st}^{ii}$ with $x_{st}^{ii} \geq 0$ and $\bar{x}_{st}^{ii} \geq 0$. After combining everything the full program reads as

$$E_{\text{DC-Linear}}(\mathbf{x}, \mathbf{y}) = \sum_{s,i} (f_s^i)_{\circlearrowleft} (x_s^i, y_s^i) + \sum_{s \sim t} \sum_i \sum_{k \in \{0, \pm 1\}} (f_{st}^{i,i+k})_{\circlearrowleft} (x_{st}^{i,i+k}, y_{st}^{i,i+k}) \tag{14}$$

$$\begin{aligned} \text{s.t. } x_s^i &= x_{st}^{ii} + x_{st}^{i,i+1} + x_{st}^{i,i-1} - \bar{x}_{st}^{ii}, & x_t^i &= x_{st}^{ii} + x_{st}^{i-1,i} + x_{st}^{i+1,i} - \bar{x}_{st}^{ii}, \\ y_s^i &= y_{st \rightarrow s}^{ii} + y_{st \rightarrow s}^{i,i+1} + y_{st \rightarrow s}^{i,i-1}, & y_t^i &= y_{st \rightarrow t}^{ii} + y_{st \rightarrow t}^{i-1,i} + y_{st \rightarrow t}^{i+1,i}, \\ x_s &\in \Delta^L, y_s^i \in [0, x_s^i], x_{st}^{i,i+\{0,\pm 1\}} \geq 0, y_{st}^{i,i+\{0,\pm 1\}} \in [0, x_{st}^{i,i+\{0,\pm 1\}}]. \end{aligned}$$

Unlike $E_{\text{LP-Linear}}$, which is equivalent to $E_{\text{LP-MRF-}L^1}$, we do not yet know whether $E_{\text{DC-Linear}}$ and $E_{\text{DC-MRF}}$ are equivalent for L^1 smoothness. We observed in our experiments, that minimizers of $E_{\text{DC-Linear}}$ have a slightly higher (true) energy than solutions of $E_{\text{DC-MRF}}$, but the numerical results are not fully conclusive.

5 Numerical Results

In our numerical experiments we compare several methods: the baseline method is the LP relaxation with a natural number of labels (256 for grayscale image processing, 64 for dense stereo computation). Since we have L^1 smoothness priors, we solve a respective minimum cut problem to find the minimizer. We call the result “baseline” rather than ground truth, since the continuous label space is still discretized. In order to judge the benefits of using convex potentials over constant ones, we also apply the minimum cut on compressed sets of labels, where the unary potentials are averaged (“GC-Avg”) or subsampled versions of the original ones (“GC-Sampled”). Further, the impact of replacing the original unary potentials by convex segments is demonstrated by running the baseline method using convexified unary costs (“GC-Convex”). The full general model described in Section 4.3 is denoted as “DC-MRF”, and the construction for L^1 smoothness priors is depicted as “DC-Linear” in the following. Comparing the results of DC-MRF/DC-Linear and the ones of GC-Convex indicate the tightness of our proposed convex models.




Fig. 3. Fitting piecewise convex quadratic unary potentials (a) and lower envelopes (b) to given data costs

We implemented prototype algorithms for finding minimizers of $E_{\text{DC-MRF}}$ for L^1 smoothness priors and $E_{\text{DC-Linear}}$ using the simultaneous-direction method of multipliers (SDMM, see e.g. [26] for a comprehensive overview of proximal methods for non-smooth convex optimization). General unary potentials are transformed into a convex surrogate in each label range by fitting a quadratic polynomial, such that the unary costs at the end-points are preserved. If the fitted polynomial is concave, a linear approximation is used. Fig. 3 depicts two cost profiles, a truncated quadratic data term of a robust ROF model and the matching costs used for stereo, and the corresponding piecewise convex surrogates.

5.1 Robust ROF Energy

In order to determine the influence of reducing the number of ranges in the label space, we utilize a robust version of the ROF energy, that replaces the quadratic data term by the non-convex truncated one,

$$E_{\text{robust-ROF}}(u; f) = \frac{\lambda}{2} \sum_s \min \{ (u_s - f_s)^2, T \} + \sum_{s \sim t} |u_s - u_t|. \quad (15)$$

Our choice of T is 500, hence deviations of about 22 grayscale units are considered as outliers. λ is set to $1/5$. Figure 4(a) illustrates the baseline (reference) result, and the minimizer of $E_{\text{DC-MRF}}$ using 8 ranges for the label space $[0, 255]$ is shown in Fig. 4(b). Further, a comparison in terms of PSNR between the baseline and the results of the tested methods with varying numbers of ranges is displayed in Fig. 4(c). From the graph we can learn two things: first, the discrete-continuous models DC-MRF and DC-Linear are very close to the 256-label result with convexified data term. Second, DC-MRF and DC-Linear are far superior to naively reducing the number of discrete labels as in GC-Avg and GC-Sampled. A full visual illustration of the obtained results and corresponding values of $E_{\text{robust-ROF}}$ are provided in the supplementary material. The supplementary material also illustrates the evolution of energies with respect to CPU time for the baseline approach and for the proposed discrete-continuous formulation. In a nutshell minimizing $E_{\text{DC-Linear}}$ quickly reaches an energy close to the converged value for the robust ROF problem.




Fig. 4. Visual comparison between the baseline solution (a) and the DC-MRF approach (b) for the robust ROF denoising model. (c) illustrates the image differences to the baseline solution for different methods. We were not able to minimize the full E_{DC-MRF} for 32 subranges due to memory consumption.




Fig. 5. Visual comparison for computational stereo using an L^1 smoothness prior. The columns correspond to numbers of used segment to represent the range $[0, 128]$. The rightmost column always shows the (globally optimal) result of a max-flow approach using 128 discrete labels for reference. 1st row: max-flow with sampled unary potentials. 2nd row: max-flow with averaged unary potentials. 3rd row: full 128-label maxflow, but with convexified unary potentials in each subrange. 4th row: minimizers of $E_{DC-Linear}$. The differences between the 3rd and the 4th row illustrate how much a convexified unary potential affects the minimizer and what is lost due to the compressed discrete-continuous representation in $E_{DC-Linear}$.

5.2 Dense Stereo

Since the unary potentials of the robust ROF energy are very structured, we illustrate our discrete-continuous MRF approach for the non-parametric data term induced by stereo matching costs. We use the truncated L^1 difference between 3×3 patches after Sobel-filtering the (grayscale) input images, since this is a realistic similarity measure for stereo. The truncation threshold is set to 20. Since the full energy $E_{\text{DC-MRF}}$ is too expensive to optimize, we report visual results for DC-Linear in Fig. 5. Following [8] we use a non-parametric convex envelope to obtain a piece-wise convex approximation to the original matching costs. The supplementary material also shows results for a parametric convex fit using piece-wise quadratic surrogates.

6 Summary and Future Work

We presented a novel, discrete-continuous formulation to solve Markov random fields over continuous label spaces. The proposed energy model is a proper extension of widely-used relaxations for labeling problems. While the proposed convex approach is not a perfectly tight relaxation in general, preliminary numerical results are promising. There are several directions for future work: most importantly, finding more efficient algorithms than our prototype implementation to minimize the DC-MRF energy will enable applications for larger-scale problems. Further, an improved theoretical analysis of the construction described in Section 4.4 and respective generalizations are of high relevance. Finally, avoiding the grid bias in the solution and using a more isotropic regularization approach are beneficial in low-level vision applications and subject of future research.

References

1. Boykov, Y., Jolly, M.P.: Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images. In: Proc. ICCV, pp. 105–112 (2001)
2. Beale, E.M.L., Forrest, J.J.H.: Global optimization using special ordered sets. *Mathematical Programming* 10(1), 52–69 (1976)
3. Wainwright, M.J., Jordan, M.I.: Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.* 1, 1–305 (2008)
4. Werner, T.: A linear programming approach to max-sum problem: A review. *IEEE Trans. Pattern Anal. Mach. Intell.* 29(7) (2007)
5. Ishikawa, H., Geiger, D.: Occlusions, Discontinuities, and Epipolar Lines in Stereo. In: Burkhardt, H.-J., Neumann, B. (eds.) *ECCV 1998*. LNCS, vol. 1406, pp. 232–248. Springer, Heidelberg (1998)
6. Boykov, Y., Veksler, O., Zabih, R.: Markov random fields with efficient approximations. In: Proc. CVPR, pp. 648–655 (1998)
7. Ishikawa, H.: Exact optimization for Markov random fields with convex priors. *IEEE Trans. Pattern Anal. Mach. Intell.* 25(10), 1333–1336 (2003)
8. Bhusnurmath, A., Taylor, C.J.: Stereo matching using interior point methods. In: Proc. 3DPVT (2008)

9. Lucas, B., Kanade, T.: An iterative image registration technique with an application to stereo vision. In: *Int. Joint Conf. on Artificial Intell.*, pp. 674–679 (1981)
10. Horn, B.K.P., Schunck, B.G.: Determining optical flow. *Artificial Intelligence* 17, 185–203 (1981)
11. Rother, C., Kohli, P., Feng, W., Jia, J.: Minimizing sparse higher order energy functions of discrete variables. In: *Proc. CVPR* (2009)
12. Kohli, P., Kumar, M.P.: Energy minimization for linear envelope MRFs. In: *Proc. CVPR* (2010)
13. Komodakis, N., Paragios, N.: Beyond pairwise energies: Efficient optimization for higher-order MRFs. In: *Proc. CVPR* (2009)
14. Alberti, G., Bouchitté, G., Maso, G.D.: The calibration method for the Mumford-Shah functional and free-discontinuity problems. *Calc. Var. Partial Differential Equations* 16(3), 299–333 (2003)
15. Chambolle, A., Cremers, D., Pock, T.: A convex approach for computing minimal partitions. Technical report, Ecole Polytechnique (2008)
16. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: An algorithm for minimizing the piecewise smooth Mumford-Shah functional. In: *Proc. ICCV* (2009)
17. Strekalovskiy, E., Goldluecke, B., Cremers, D.: Tight convex relaxations for vector-valued labeling problems. In: *Proc. ICCV* (2011)
18. Peng, J., Hazan, T., McAllester, D., Urtasun, R.: Convex max-product algorithms for continuous MRFs with applications to protein folding. In: *Proc. ICML* (2011)
19. Glocker, B., Paragios, N., Komodakis, N., Tziritas, G., Navab, N.: Optical flow estimation with uncertainties through dynamic MRFs. In: *Proc. CVPR* (2008)
20. Kolmogorov, V.: Convergent tree-reweighted message passing for energy minimization. *IEEE Trans. Pattern Anal. Mach. Intell.* 28(10), 1568–1583 (2006)
21. Globerson, A., Jaakkola, T.: Fixing max-product: Convergent message passing algorithms for MAP LP-relaxations. In: *NIPS* (2007)
22. Weiss, Y., Yanover, C., Meltzer, T.: MAP estimation, linear programming and belief propagation with convex free energies. In: *UAI* (2007)
23. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D* 60, 259–268 (1992)
24. Chambolle, A.: An algorithm for total variation minimization and applications. *J. Math. Imag. Vision* 20(1-2), 89–97 (2004)
25. Zach, C., Häne, C., Pollefeys, M.: What is optimized in convex relaxations for multi-label problems: Connecting discrete and continuously-inspired MAP inference. Technical report, Microsoft Research (2012)
26. Combettes, P.L., Pesquet, J.C.: Proximal Splitting Methods in Signal Processing. In: *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pp. 185–212. Springer (2011)