

Sublabel-Accurate Convex Relaxation of Vectorial Multilabel Energies

Emanuel Laude^{*1}, Thomas Möllenhoff^{*1}, Michael Moeller¹,
Jan Lellmann², and Daniel Cremers¹

¹Technical University of Munich^{**}

²University of Lübeck

Abstract. Convex relaxations of multilabel problems have been demonstrated to produce provably optimal or near-optimal solutions to a variety of computer vision problems. Yet, they are of limited practical use as they require a fine discretization of the label space, entailing a huge demand in memory and runtime. In this work, we propose the first sub-label accurate convex relaxation for vectorial multilabel problems. Our key idea is to approximate the dataterm in a piecewise convex (rather than piecewise linear) manner. As a result we have a more faithful approximation of the original cost function that provides a meaningful interpretation for fractional solutions of the relaxed convex problem.

Keywords: Convex Relaxation, Optimization, Variational Methods




Fig. 1: In (a) we show a nonconvex dataterm. Convexification without lifting would result in the energy (b). Classical lifting methods [11] (c), approximate the energy piecewise linearly between the labels, whereas the proposed method results in an approximation that is convex on each triangle (d). Therefore, we are able to capture the structure of the nonconvex energy much more accurately.

^{*} These authors contributed equally.

^{**} This work was supported by the ERC Starting Grant “Convex Vision”.

1 Introduction

1.1 Nonconvex Vectorial Problems

In this paper, we derive a sublabel-accurate convex relaxation for vectorial optimization problems of the form

$$\min_{u: \Omega \rightarrow \Gamma} \int_{\Omega} \rho(x, u(x)) \, dx + \lambda TV(u), \quad (1)$$

where $\Omega \subset \mathbb{R}^d$, $\Gamma \subset \mathbb{R}^n$ and $\rho : \Omega \times \Gamma \rightarrow \mathbb{R}$ denotes a generally nonconvex pointwise dataterm. As regularization we focus on the *total variation* defined as:

$$TV(u) = \sup_{q \in C_c^\infty(\Omega, \mathbb{R}^{n \times d}), \|q(x)\|_{S^\infty} \leq 1} \int_{\Omega} \langle u, \text{Div } q \rangle \, dx, \quad (2)$$

where $\|\cdot\|_{S^\infty}$ is the Schatten- ∞ norm on $\mathbb{R}^{n \times d}$, i.e., the largest singular value. For differentiable functions u we can integrate (2) by parts to find

$$TV(u) = \int_{\Omega} \|\nabla u(x)\|_{S^1} \, dx, \quad (3)$$

where the dual norm $\|\cdot\|_{S^1}$ penalizes the sum of the singular values of the Jacobian, which encourages the individual components of u to jump in the same direction. This type of regularization is part of the framework of Sapiro and Ringach [19].

1.2 Related Work

Due to its nonconvexity the optimization of (1) is challenging. For the scalar case ($n = 1$), Ishikawa [9] proposed a pioneering technique to obtain globally optimal solutions in a spatially discrete setting, given by the minimum s-t-cut of a graph representing the space $\Omega \times \Gamma$. A continuous formulation was introduced by Pock et al. [15] exhibiting several advantages such as less grid bias and parallelizability.

In a series of papers [16,14], connections of the above approaches were made to the mathematical theory of *cartesian currents* [6] and the calibration method for the Mumford-Shah functional [1], leading to a generalization of the convex relaxation framework [15] to more general (in particular nonconvex) regularizers.

In the following, researchers have strived to generalize the concept of functional lifting and convex relaxation to the vectorial setting ($n > 1$). If the dataterm and the regularizer are both separable in the label dimension, one can simply apply the above convex relaxation approach in a channel-wise manner

to each component separately. But when either the dataterm or the regularizer couple the label components, the situation becomes more complex [8,20].

The approach which is most closely related to our work, and which we consider as a baseline method, is the one by Lellmann et al. [11]. They consider coupled dataterms with coupled total variation regularization of the form (2).

A drawback shared by all mentioned papers is that ultimately one has to discretize the label space. While Lellmann et al. [11] propose a sublabel-accurate regularizer, we show that their dataterm leads to solutions which still have a strong bias towards the label grid. For the scalar-valued setting, continuous label spaces have been considered in the MRF community by Zach et al. [22] and Fix et al. [5]. The paper [21] proposes a method for mixed continuous and discrete vectorial label spaces, where everything is derived in the spatially discrete MRF setting. Möllenhoff et al. [12] recently proposed a novel formulation of the scalar-valued case which retains fully continuous label spaces even after discretization. The contribution of this work is to extend [12] to vectorial label spaces, thereby complementing [11] with a sublabel-accurate dataterm.

1.3 Contribution

In this work we propose the first sublabel-accurate convex formulation of vectorial labeling problems. It generalizes the formulation for scalar-valued labeling problems [12] and thus includes important applications such as optical flow estimation or color image denoising. We show that our method, derived in a spatially continuous setting, has a variety of interesting theoretical properties as well as practical advantages over the existing labeling approaches:

- We generalize existing functional lifting approaches (see Sec. 2.2).
- We show that our method is the best convex under-approximation (in a local sense), see Prop. 1 and Prop. 2.
- Due to its sublabel-accuracy our method requires only a small amount of labels to produce good results which leads to a drastic reduction in memory. We believe that this is a vital step towards the real-time capability of lifting and convex relaxation methods. Moreover, our method eliminates the label bias, that previous lifting methods suffer from, even for many labels.
- In Sec. 2.3 we propose a regularizer that couples the different label components by enforcing a joint jump normal. This is in contrast to [8], where the components are regularized separately.
- For convex dataterms, our method is equivalent to the unlifted problem – see Prop. 4. Therefore, it allows a seamless transition between direct optimization and convex relaxation approaches.

1.4 Notation

We write $\langle x, y \rangle = \sum_i x_i y_i$ for the standard inner product on \mathbb{R}^n or the Frobenius product if x, y are matrices. Similarly $\|\cdot\|$ without any subscript denotes the usual Euclidean norm, respectively the Frobenius norm for matrices.

We denote the convex conjugate of a function $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ by $f^*(y) = \sup_{x \in \mathbb{R}^n} \langle y, x \rangle - f(x)$. It is an important tool for devising convex relaxations, as the biconjugate f^{**} is the largest lower-semicontinuous (lsc.) convex function below f . For the indicator function of a set C we write δ_C , i.e., $\delta_C(x) = 0$ if $x \in C$ and ∞ otherwise. $\Delta_n^U \subset \mathbb{R}^n$ stands for the unit n -simplex.

2 Convex Formulation

2.1 Lifted Representation

Motivated by Fig. 1, we construct an equivalent representation of (1) in a higher dimensional space, before taking the convex envelope.

Let $\Gamma \subset \mathbb{R}^n$ be a compact and convex set. We partition Γ into a set \mathcal{T} of n -simplices Δ_i so that Γ is a disjoint union of Δ_i up to a set of measure zero. Let t^{ij} be the j -th vertex of Δ_i and denote by $\mathcal{V} = \{t^1, \dots, t^{|\mathcal{V}|}\}$ the union of all vertices, referred to as labels, with $1 \leq i \leq |\mathcal{T}|$, $1 \leq j \leq n+1$ and $1 \leq i_j \leq |\mathcal{V}|$. For $u : \Omega \rightarrow \Gamma$, we refer to $u(x)$ as a *sublabel*. Any sublabel can be written as a convex combination of the vertices of a simplex Δ_i with $1 \leq i \leq |\mathcal{T}|$ for appropriate barycentric coordinates $\alpha \in \Delta_n^U$:

$$u(x) = T_i \alpha := \sum_{j=1}^{n+1} \alpha_j t^{ij}, \quad T_i := (t^{i1}, t^{i2}, \dots, t^{i(n+1)}) \in \mathbb{R}^{n \times (n+1)}. \quad (4)$$

By encoding the vertices $t^k \in \mathcal{V}$ using a one-of- $|\mathcal{V}|$ representation e^k we can identify any $u(x) \in \Gamma$ with a sparse vector $\mathbf{u}(x)$ containing at least $|\mathcal{V}| - n$ many zeros and vice versa:

$$\begin{aligned} \mathbf{u}(x) &= E_i \alpha := \sum_{j=1}^{n+1} \alpha_j e^{ij}, \quad E_i := (e^{i1}, e^{i2}, \dots, e^{i(n+1)}) \in \mathbb{R}^{|\mathcal{V}| \times (n+1)}, \\ \mathbf{u}(x) &= \sum_{k=1}^{|\mathcal{V}|} t^k \mathbf{u}_k(x), \quad \alpha \in \Delta_n^U, \quad 1 \leq i \leq |\mathcal{T}|. \end{aligned} \quad (5)$$

The entries of the vector e^{ij} are zero except for the (i_j) -th entry, which is equal to one. We refer to $\mathbf{u} : \Omega \rightarrow \mathbb{R}^{|\mathcal{V}|}$ as the *lifted* representation of u . This one-to-one-correspondence between $u(x) = T_i \alpha$ and $\mathbf{u}(x) = E_i \alpha$ is shown in Fig. 2. Note that both, α and i depend on x . However, for notational convenience we drop the dependence on x whenever we consider a fixed point $x \in \Omega$.




Fig. 2: This figure illustrates our notation and the one-to-one correspondence between $u(x) = (0.3, 0.2)^\top$ and the lifted $\mathbf{u}(x)$ containing the barycentric coordinates $\alpha = (0.7, 0.1, 0.2)^\top$ of the sublabel $u(x) \in \Delta_4 = \text{conv}\{t^2, t^3, t^6\}$. The triangulation $(\mathcal{V}, \mathcal{T})$ of $\Gamma = [-1; 1] \times [0; 1]$ is visualized via the gray lines, corresponding to the triangles and the gray dots, corresponding to the vertices $\mathcal{V} = \{(-1, 0)^\top, (0, 0)^\top, \dots, (1, 1)^\top\}$, that we refer to as the labels.

2.2 Convexifying the Dataterm

Let for now the weight of the regularizer in (1) be zero. Then, at each point $x \in \Omega$ we minimize a generally nonconvex energy over a compact set $\Gamma \subset \mathbb{R}^n$:

$$\min_{u \in \Gamma} \rho(u). \quad (6)$$

We set up the lifted energy so that it attains finite values if and only if the argument \mathbf{u} is a sparse representation $\mathbf{u} = E_i \alpha$ of a sublabel $u \in \Gamma$:

$$\rho(\mathbf{u}) = \min_{1 \leq i \leq |\mathcal{T}|} \rho_i(\mathbf{u}), \quad \rho_i(\mathbf{u}) = \begin{cases} \rho(T_i \alpha), & \text{if } \mathbf{u} = E_i \alpha, \alpha \in \Delta_n^U, \\ \infty, & \text{otherwise.} \end{cases} \quad (7)$$

Problems (6) and (7) are equivalent due to the one-to-one correspondence of $u = T_i \alpha$ and $\mathbf{u} = E_i \alpha$. However, energy (7) is finite on a nonconvex set only. In order to make optimization tractable, we minimize its convex envelope.

Proposition 1 *The convex envelope of (7) is given as:*

$$\begin{aligned} \rho^{**}(\mathbf{u}) &= \sup_{\mathbf{v} \in \mathbb{R}^{|\mathcal{V}|}} \langle \mathbf{u}, \mathbf{v} \rangle - \max_{1 \leq i \leq |\mathcal{T}|} \rho_i^*(\mathbf{v}), \\ \rho_i^*(\mathbf{v}) &= \langle E_i b_i, \mathbf{v} \rangle + \rho_i^*(A_i^\top E_i^\top \mathbf{v}), \quad \rho_i := \rho + \delta_{\Delta_i}. \end{aligned} \quad (8)$$

b_i and A_i are given as $b_i := M_i^{n+1}$, $A_i := (M_i^1, M_i^2, \dots, M_i^n)$, where M_i^j are the columns of the matrix $M_i := (T_i^\top, \mathbf{1})^{-\top} \in \mathbb{R}^{n+1 \times n+1}$.

Proof. Follows from a calculation starting at the definition of ρ^{**} . See supplementary material for a detailed derivation.




Fig. 3: Geometrical intuition for the proposed lifting and standard lifting [11] for the special case of 1-dimensional range $\Gamma = [a, b]$ and 3 labels $\{t^1, t^2, t^3\}$. The standard lifting corresponds to a linear interpolation of the original cost in between the locations t^1, t^2, t^3 , which are associated with the vertices e^1, e^2, e^3 in the lifted energy (lower left). The proposed method extends the cost to the relaxed set in a more precise way: The original cost is preserved on the connecting lines between adjacent e^i (black lines on the bottom right) up to concave parts (red graphs and lower surface on the right). This information, which may influence the exact location of the minimizer, is lost in the standard formulation. If the solution of the lifted formulation \mathbf{u} is in the interior (gray area) an approximate solution to the original problem can still be obtained via Eq. (5).

The geometric intuition of this construction is depicted in Fig. 3. Note that if one prescribes the value of ρ_i in (7) only on the *vertices* of the unit simplices Δ_n^U , i.e., $\rho(\mathbf{u}) = \rho(t^k)$ if $\mathbf{u} = e^k$ and $+\infty$ otherwise, one obtains the linear biconjugate $\rho^{**}(\mathbf{u}) = \langle \mathbf{u}, \mathbf{s} \rangle$, $\mathbf{s} = (\rho(t^1), \dots, \rho(t^L))$ on the feasible set. This coincides with the standard relaxation of the dataterm used in [16,10,4,11]. In that sense, our approach can be seen as a relaxing the dataterm in a more precise way, by incorporating the true value of ρ not only on the finite set of labels \mathcal{V} , but also everywhere in between, i.e., on every *sublabel*.

2.3 Lifting the Vectorial Total Variation

We define the lifted vectorial total variation as

$$\mathbf{TV}(\mathbf{u}) = \int_{\Omega} \Psi(D\mathbf{u}), \quad (9)$$

where $D\mathbf{u}$ denotes the distributional derivative of \mathbf{u} and Ψ is positively one-homogeneous, i.e., $\Psi(c\mathbf{u}) = c\Psi(\mathbf{u})$, $c \geq 0$. For such functions, the meaning of (9) can be made fully precise using the polar decomposition of the Radon measure $D\mathbf{u}$ [2, Cor. 1.29, Thm. 2.38]. However, in the following we restrict ourselves to an intuitive motivation for the derivation of Ψ for smooth functions.

Our goal is to find Ψ so that $\mathbf{TV}(\mathbf{u}) = TV(u)$ whenever $\mathbf{u} : \Omega \rightarrow \mathbb{R}^{|\mathcal{V}|}$ corresponds to some $u : \Omega \rightarrow \Gamma$, in the sense that $\mathbf{u}(x) = E_i\alpha$ whenever $u(x) = T_i\alpha$. In order for the equality to hold, it must in particular hold for all u that are classically differentiable, i.e., $Du = \nabla u$, and whose Jacobian $\nabla u(x)$ is of rank 1, i.e., $\nabla u(x) = (T_i\alpha - T_j\beta) \otimes \nu(x)$ for some $\nu(x) \in \mathbb{R}^d$. This rank 1 constraint enforces the different components of u to have the same jump normal, which is desirable in many applications. In that case, we observe

$$TV(u) = \int_{\Omega} \|T_i\alpha - T_j\beta\| \cdot \|\nu(x)\| dx. \quad (10)$$

For the corresponding lifted representation \mathbf{u} , we have $\nabla\mathbf{u}(x) = (E_i\alpha - E_j\beta) \otimes \nu(x)$. Therefore it is natural to require $\Psi(\nabla\mathbf{u}(x)) = \Psi((E_i\alpha - E_j\beta) \otimes \nu(x)) := \|T_i\alpha - T_j\beta\| \cdot \|\nu(x)\|$ in order to achieve the goal $\mathbf{TV}(\mathbf{u}) = TV(u)$. Motivated by these observations, we define

$$\Psi(\mathbf{p}) := \begin{cases} \|T_i\alpha - T_j\beta\| \cdot \|\nu\| & \text{if } \mathbf{p} = (E_i\alpha - E_j\beta) \otimes \nu, \\ \infty & \text{otherwise,} \end{cases} \quad (11)$$

where $\alpha, \beta \in \Delta_{n+1}^U$, $\nu \in \mathbb{R}^d$ and $1 \leq i, j \leq |\mathcal{T}|$. Since the convex envelope of (9) is intractable, we derive a ‘‘locally’’ tight convex underapproximation:

$$\mathbf{R}(\mathbf{u}) = \sup_{\mathbf{q}: \Omega \rightarrow \mathbb{R}^{d \times |\mathcal{V}|}} \int_{\Omega} \langle \mathbf{u}, \text{Div } \mathbf{q} \rangle - \Psi^*(\mathbf{q}) dx. \quad (12)$$

Proposition 2 *The convex conjugate of Ψ is*

$$\Psi^*(\mathbf{q}) = \delta_{\mathcal{K}}(\mathbf{q}) \quad (13)$$

with convex set

$$\mathcal{K} = \bigcap_{1 \leq i, j \leq |\mathcal{T}|} \left\{ \mathbf{q} \in \mathbb{R}^{d \times |\mathcal{V}|} \mid \|Q_i\alpha - Q_j\beta\| \leq \|T_i\alpha - T_j\beta\|, \alpha, \beta \in \Delta_{n+1}^U \right\}, \quad (14)$$

and $Q_i = (\mathbf{q}^{i1}, \mathbf{q}^{i2}, \dots, \mathbf{q}^{in+1}) \in \mathbb{R}^{d \times n+1}$. $\mathbf{q}^j \in \mathbb{R}^d$ are the columns of \mathbf{q} .

Proof. Follows from a calculation starting at the definition of the convex conjugate Ψ^* . See supplementary material.

Interestingly, although in its original formulation (14) the set \mathcal{K} has infinitely many constraints, one can equivalently represent \mathcal{K} by finitely many.

Proposition 3 *The set \mathcal{K} in equation (14) is the same as*

$$\mathcal{K} = \left\{ \mathbf{q} \in \mathbb{R}^{d \times |\mathcal{V}|} \mid \|D_{\mathbf{q}}^i\|_{S^\infty} \leq 1, 1 \leq i \leq |\mathcal{T}| \right\}, \quad D_{\mathbf{q}}^i = Q_i D (T_i D)^{-1}, \quad (15)$$

where the matrices $Q_i D \in \mathbb{R}^{d \times n}$ and $T_i D \in \mathbb{R}^{n \times n}$ are given as

$$Q_i D := (\mathbf{q}^{i_1} - \mathbf{q}^{i_{n+1}}, \dots, \mathbf{q}^{i_n} - \mathbf{q}^{i_{n+1}}), \quad T_i D := (t^{i_1} - t^{i_{n+1}}, \dots, t^{i_n} - t^{i_{n+1}}).$$

Proof. Similar to the analysis in [11], equation (14) basically states the Lipschitz continuity of a piecewise linear function defined by the matrices $\mathbf{q} \in \mathbb{R}^{d \times |\mathcal{V}|}$. Therefore, one can expect that the Lipschitz constraint is equivalent to a bound on the derivative. For the complete proof, see supplementary material.

2.4 Lifting the Overall Optimization Problem

Combining dataterm and regularizer, the overall optimization problem is given

$$\min_{\mathbf{u}: \Omega \rightarrow \mathbb{R}^{|\mathcal{V}|}} \sup_{\mathbf{q}: \Omega \rightarrow \mathcal{K}} \int_{\Omega} \rho^{**}(\mathbf{u}) + \langle \mathbf{u}, \text{Div } \mathbf{q} \rangle \, dx. \quad (16)$$

A highly desirable property is that, opposed to any other vectorial lifting approach from the literature, our method with just one simplex applied to a convex problem yields the same solution as the unlifted problem.

Proposition 4 *If the triangulation contains only 1 simplex, $\mathcal{T} = \{\Delta\}$, i.e., $|\mathcal{V}| = n + 1$, then the proposed optimization problem (16) is equivalent to*

$$\min_{u: \Omega \rightarrow \Delta} \int_{\Omega} (\rho + \delta_{\Delta})^{**}(x, u(x)) \, dx + \lambda TV(u), \quad (17)$$

which is (1) with a globally convexified dataterm on Δ .

Proof. For $u = t^{n+1} + TD\tilde{u}$ the substitution $\mathbf{u} = (\tilde{u}_1, \dots, \tilde{u}_n, 1 - \sum_{j=1}^n \tilde{u}_j)$ into ρ^{**} and \mathbf{R} yields the result. For a complete proof, see supplementary material.

3 Numerical Optimization

3.1 Discretization

For now assume that $\Omega \subset \mathbb{R}^d$ is a d -dimensional Cartesian grid and let Div denote a finite-difference divergence operator with $\text{Div } \mathbf{q} : \Omega \rightarrow \mathbb{R}^{|\mathcal{V}|}$. Then the relaxed energy minimization problem becomes

$$\min_{\mathbf{u}: \Omega \rightarrow \mathbb{R}^{|\mathcal{V}|}} \max_{\mathbf{q}: \Omega \rightarrow \mathcal{K}} \sum_{x \in \Omega} \rho^{**}(x, \mathbf{u}(x)) + \langle \text{Div } \mathbf{q}, \mathbf{u} \rangle. \quad (18)$$

In order to get rid of the pointwise maximum over $\rho_i^*(\mathbf{v})$ in Eq. (8), we introduce additional variables $w(x) \in \mathbb{R}$ and additional constraints $(\mathbf{v}(x), w(x)) \in \mathcal{C}$, $x \in \Omega$ so that $w(x)$ attains the value of the pointwise maximum:

$$\min_{\mathbf{u}: \Omega \rightarrow \mathbb{R}^{|\mathcal{V}|}} \max_{\substack{(\mathbf{v}, w): \Omega \rightarrow \mathcal{C} \\ \mathbf{q}: \Omega \rightarrow \mathcal{K}}} \sum_{x \in \Omega} \langle \mathbf{u}(x), \mathbf{v}(x) \rangle - w(x) + \langle \text{Div } \mathbf{q}, \mathbf{u} \rangle, \quad (19)$$

where the set \mathcal{C} is given as

$$\mathcal{C} = \bigcap_{1 \leq i \leq |\mathcal{T}|} \mathcal{C}_i, \quad \mathcal{C}_i := \left\{ (x, y) \in \mathbb{R}^{|\mathcal{V}|+1} \mid \rho_i^*(x) \leq y \right\}. \quad (20)$$

For numerical optimization we use a GPU-based implementation¹ of a first-order primal-dual method [14]. The algorithm requires the orthogonal projections of the dual variables onto the sets \mathcal{C} respectively \mathcal{K} in every iteration. However, the projection onto an epigraph of dimension $|\mathcal{V}| + 1$ is difficult for large values of $|\mathcal{V}|$. We rewrite the constraints $(\mathbf{v}(x), w(x)) \in \mathcal{C}_i$, $1 \leq i \leq |\mathcal{T}|$, $x \in \Omega$ as $(n + 1)$ -dimensional epigraph constraints introducing variables $r^i(x) \in \mathbb{R}^n$, $s_i(x) \in \mathbb{R}$:

$$\rho_i^*(r^i(x)) \leq s_i(x), \quad r^i(x) = A_i^\top E_i^\top \mathbf{v}(x), \quad s_i(x) = w(x) - \langle E_i \mathbf{b}_i, \mathbf{v}(x) \rangle. \quad (21)$$

These equality constraints can be implemented using Lagrange multipliers. For the projection onto the set \mathcal{K} we use an approach similar to [7, Figure 7].

3.2 Epigraphical Projections

Computing the Euclidean projection onto the epigraph of ρ_i^* is a central part of the numerical implementation of the presented method. However, for $n > 1$ this is nontrivial. Therefore we provide a detailed explanation of the projection methods used for different classes of ρ_i . We will consider quadratic, truncated quadratic and piecewise linear ρ .

¹ https://github.com/tum-vision/sublabel_relax

Quadratic case: Let ρ be of the form $\rho(u) = \frac{a}{2} u^\top u + b^\top u + c$. A direct projection onto the epigraph of $\rho_i^* = (\rho + \delta_{\Delta_i})^*$ for $n > 1$ is difficult. However, the epigraph can be decomposed into separate epigraphs for which it is easier to project onto: For proper, convex, lsc. functions f, g the epigraph of $(f + g)^*$ is the Minkowski sum of the epigraphs of f^* and g^* (cf. [17, Exercise 1.28, Theorem 11.23a]). This means that it suffices to compute the projections onto the epigraphs of a quadratic function $f^* = \rho^*$ and a convex, piecewise linear function $g^*(v) = \max_{1 \leq j \leq n+1} \langle t^{ij}, v \rangle$ by rewriting constraint (21) as

$$\rho^*(r_f) \leq s_f, \delta_{\Delta_i}^*(c_g) \leq d_g \text{ s.t. } (r, s) = (r_f, s_f) + (c_g, d_g). \quad (22)$$

For the projection onto the epigraph of a n -dimensional quadratic function we use the method described in [20, Appendix B.2]. The projection onto a piecewise linear function is described in the last paragraph of this section.

Truncated quadratic case: Let ρ be of the form $\rho(u) = \min \{ \nu, \frac{a}{2} u^\top u + b^\top u + c \}$ as it is the case for the nonconvex robust ROF with a truncated quadratic dataterm in Sec. 4.2. Again, a direct projection onto the epigraph of ρ_i^* is difficult. However, a decomposition of the epigraph into simpler epigraphs is possible as the epigraph of $\min\{f, g\}^*$ is the intersection of the epigraphs of f^* and g^* . Hence, one can separately project onto the epigraphs of $(\nu + \delta_{\Delta_i})^*$ and $(\frac{a}{2} u^\top u + b^\top u + c + \delta_{\Delta_i})^*$. Both of these projections can be handled using the methods from the other paragraphs.

Piecewise linear case: In case ρ is piecewise linear on each Δ_i , i.e., ρ attains finite values at a discrete set of sampled sublabeled $\mathcal{V}_i \subset \Delta_i$ and interpolates linearly between them, we have that

$$(\rho + \delta_{\Delta_i})^*(v) = \max_{\tau \in \mathcal{V}_i} \langle \tau, v \rangle - \rho(\tau). \quad (23)$$

Again this is a convex, piecewise linear function. For the projection onto the epigraph of such a function, a quadratic program of the form

$$\min_{(x, y) \in \mathbb{R}^{n+1}} \frac{1}{2} \|x - c\|^2 + \frac{1}{2} \|y - d\|^2 \text{ s.t. } \langle \tau, x \rangle - \rho(\tau) \leq y, \forall \tau \in \mathcal{V}_i \quad (24)$$

needs to be solved. We implemented the primal active-set method described in [13, Algorithm 16.3], and found it solves the program in a few (usually 2 – 10) iterations for a moderate number of constraints.

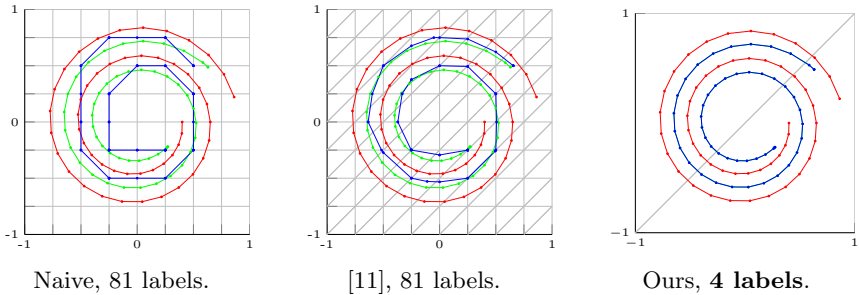


Fig. 4: ROF denoising of a vector-valued signal $f : [0, 1] \rightarrow [-1, 1]^2$, discretized on 50 points (shown in red). We compare the proposed approach (right) with two alternative techniques introduced in [11] (left and middle). The labels are visualized by the gray grid. While the naive (standard) multilabel approach from [11] (left) provides solutions that are constrained to the chosen set of labels, the sublabel accurate regularizer from [11] (middle) does allow sublabel solutions, yet – due to the dataterm bias – these still exhibit a strong preference for the grid points. In contrast, the proposed approach does not exhibit any visible grid bias providing fully sublabel-accurate solutions: With only 4 labels, the computed solutions (shown in blue) coincide with the “unlifted” problem (green).

4 Experiments

4.1 Vectorial ROF Denoising

In order to validate experimentally, that our model is exact for convex dataterms, we evaluate it on the Rudin-Osher-Fatemi [18] (ROF) model with vectorial TV (2). In our model this corresponds to defining $\rho(x, u(x)) = \frac{1}{2}\|u(x) - I(x)\|^2$. As expected based on Prop. 4 the energy of the solution of the unlifted problem is equal to the energy of the projected solution of our method for $|\mathcal{V}| = 4$ up to machine precision, as can be seen in Fig. 4 and Fig. 5. We point out, that the sole purpose of this experiment is a proof of concept as our method introduces an overhead and convex problems can be solved via direct optimization. It can be seen in Fig. 4 and Fig. 5, that the baseline method [11] has a strong label bias.

4.2 Denoising with Truncated Quadratic Dataterm

For images degraded with both, Gaussian and salt-and-pepper noise we define the dataterm as $\rho(x, u(x)) = \min \left\{ \frac{1}{2}\|u(x) - I(x)\|^2, \nu \right\}$. We solve the problem

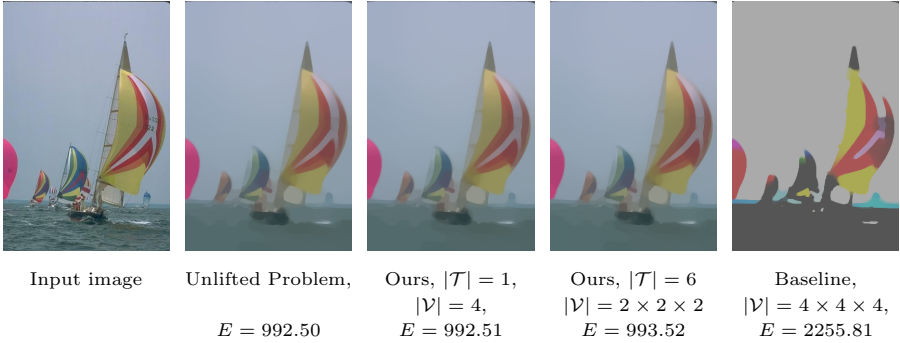


Fig. 5: Convex ROF with vectorial TV. Direct optimization and proposed method yield the same result. In contrast to the baseline method [11] the proposed approach has no discretization artefacts and yields a lower energy. The regularization parameter is chosen as $\lambda = 0.3$.




Fig. 6: ROF with a truncated quadratic dataterm ($\lambda = 0.03$ and $\nu = 0.025$). Compared to the baseline method [11] the proposed approach yields much better results, already with a very small number of 4 labels.

using the epigraph decomposition described in the second paragraph of Sec. 3.2. It can be seen, that increasing the number of labels $|\mathcal{V}|$ leads to lower energies and at the same time to a reduced effect of the TV. This occurs as we always compute a piecewise convex underapproximation of the original nonconvex dataterm, that gets tighter with a growing number of labels. The baseline method [11] again produces strong discretization artefacts even for a large number of labels $|\mathcal{V}| = 4 \times 4 \times 4 = 64$.

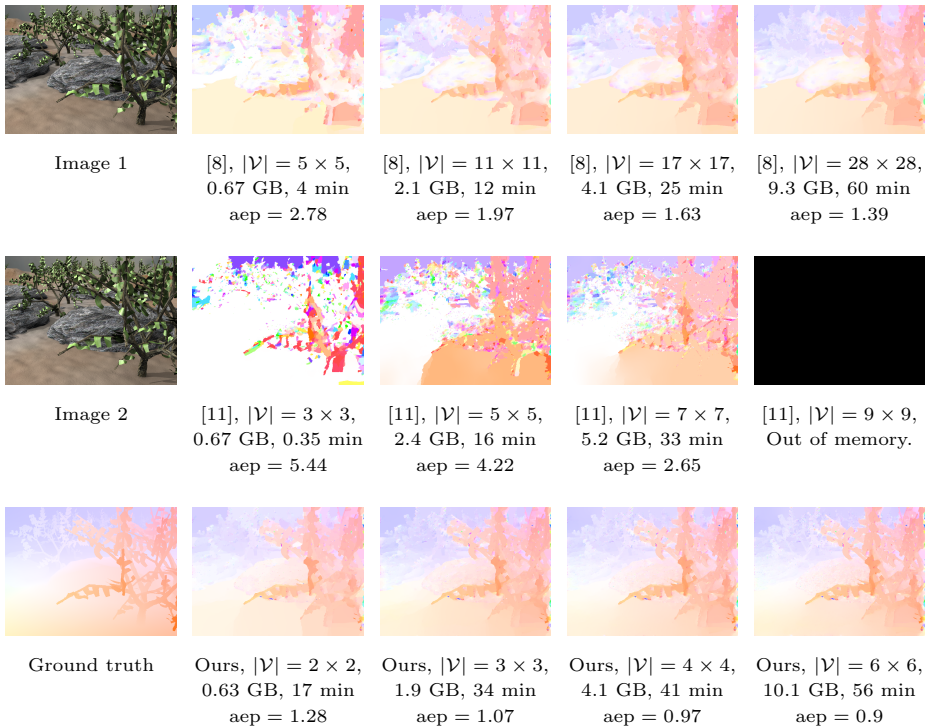


Fig. 7: We compute the optical flow using our method, the product space approach [8] and the baseline method [11] for a varying amount of labels and compare the average endpoint error (aep). The product space method clearly outperforms the baseline, but our approach finds the overall best result already with 2×2 labels. To achieve a similarly precise result as the product space method, we require 150 times fewer labels, 10 times less memory and 3 times less time. For the same number of labels, the proposed approach requires more memory as it has to store a convex approximation of the energy instead of a linear one.

4.3 Optical Flow

We compute the optical flow $v : \Omega \rightarrow \mathbb{R}^2$ between two input images I_1, I_2 . The label space $\Gamma = [-d, d]^2$ is chosen according to the estimated maximum displacement $d \in \mathbb{R}$ between the images. The dataterm is $\rho(x, v(x)) = \|I_2(x) - I_1(x + v(x))\|$, and $\lambda(x)$ is based on the norm of the image gradient $\nabla I_1(x)$.

In Fig. 7 we compare the proposed method to the product space approach [8]. Note that we implemented the product space dataterm using Lagrange mul-




Fig. 8: Large displacement flow between two 640×480 images (a) using a 81×81 search window. The result of our method with 4 labels is shown in (b), the baseline [11] in (c). Our method can correctly identify the large motion.

tipliers, also referred to as the *global* approach in [8]. While this increases the memory consumption, it comes with lower computation time and guaranteed convergence. For our method, we sample the label space $\Gamma = [-15, 15]^2$ on 150×150 sublabeled and subsequently convexify the energy on each triangle using the quickhull algorithm [3]. For the product space approach we sample the label space at equidistant labels, from 5×5 to 27×27 . As the regularizer from the product space approach is different from the proposed one, we chose μ differently for each method. For the proposed method, we set $\mu = 0.5$ and for the product space and baseline approach $\mu = 3$. We can see in Fig. 7, our method outperforms the product space approach w.r.t. the average end-point error. Our method outperforms previous lifting approaches: In Fig. 8 we compare our method on large displacement optical flow to the baseline [11]. To obtain competitive results on the Middlebury benchmark, one would need to engineer a better dataterm.

5 Conclusions

We proposed the first sublabeled-accurate convex relaxation of vectorial multilabel problems. To this end, we approximate the generally nonconvex dataterm in a piecewise convex manner as opposed to the piecewise linear approximation done in the traditional functional lifting approaches. This assures a more faithful approximation of the original cost function and provides a meaningful interpretation for the non-integral solutions of the relaxed convex problem. In experimental validations on large-displacement optical flow estimation and color image denoising, we show that the computed solutions have superior quality to the traditional convex relaxation methods while requiring substantially less memory and runtime.

References

1. Alberti, G., Bouchitté, G., Maso, G.D.: The calibration method for the Mumford-Shah functional and free-discontinuity problems. *Calc. Var. Partial Dif.* 3(16), 299–333 (2003)
2. Ambrosio, L., Fusco, N., Pallara, D.: Functions of bounded variation and free discontinuity problems. *Oxford Mathematical Monographs*, The Clarendon Press Oxford University Press, New York (2000)
3. Barber, C.B., Dobkin, D.P., Huhdanpaa, H.: The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software (TOMS)* 22(4), 469–483 (1996)
4. Chambolle, A., Cremers, D., Pock, T.: A convex approach to minimal partitions. *SIAM Journal on Imaging Sciences* 5(4), 1113–1158 (2012)
5. Fix, A., Agarwal, S.: Duality and the continuous graphical model. In: *Computer Vision ECCV 2014, Lecture Notes in Computer Science*, vol. 8691, pp. 266–281. Springer International Publishing (2014), http://dx.doi.org/10.1007/978-3-319-10578-9_18
6. Giaquinta, M., Modica, G., Souček, J.: Cartesian currents in the calculus of variations I, II, *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3., vol. 37-38*. Springer-Verlag, Berlin (1998)
7. Goldluecke, B., Strelakovski, E., Cremers, D.: The natural total variation which arises from geometric measure theory. *SIAM Journal on Imaging Sciences* 5(2), 537–563 (2012)
8. Goldluecke, B., Strelakovski, E., Cremers, D.: Tight convex relaxations for vector-valued labeling. *SIAM Journal on Imaging Sciences* 6(3), 1626–1664 (2013)
9. Ishikawa, H.: Exact optimization for Markov random fields with convex priors. *IEEE Trans. Pattern Analysis and Machine Intelligence* 25(10), 1333–1336 (2003)
10. Lellmann, J., Schnörr, C.: Continuous multiclass labeling approaches and algorithms. *SIAM Journal on Imaging Sciences* 4(4), 1049–1096 (2011)
11. Lellmann, J., Strelakovski, E., Koetter, S., Cremers, D.: Total variation regularization for functions with values in a manifold. In: *ICCV* (December 2013)
12. Möllenhoff, T., Laude, E., Moeller, M., Lellmann, J., Cremers, D.: Sublabel-accurate relaxation of nonconvex energies. In: *CVPR* (2016)
13. Nocedal, J., Wright, S.J.: *Numerical Optimization*. Springer, New York, 2nd edn. (2006)
14. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: An algorithm for minimizing the piecewise smooth Mumford-Shah functional. In: *ICCV* (2009)
15. Pock, T., Schoenemann, T., Graber, G., Bischof, H., Cremers, D.: A convex formulation of continuous multi-label problems. In: *European Conference on Computer Vision (ECCV)*. Marseille, France (October 2008)
16. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: Global solutions of variational models with convex regularization. *SIAM J. Imaging Sci.* 3(4), 1122–1145 (2010)
17. Rockafellar, R., Wets, R.B.: *Variational Analysis*. Springer (1998)

18. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* 60(1), 259–268 (1992)
19. Sapiro, G., Ringach, D.: Anisotropic diffusion of multivalued images with applications to color filtering. *IEEE Trans. Img. Proc.* 5(11), 1582–1586 (1996)
20. Strelakovsky, E., Chambolle, A., Cremers, D.: Convex relaxation of vectorial problems with coupled regularization. *SIAM Journal on Imaging Sciences* 7(1), 294–336 (2014)
21. Zach, C.: Dual decomposition for joint discrete-continuous optimization. In: *AISTATS*. pp. 632–640 (2013)
22. Zach, C., Kohli, P.: A convex discrete-continuous approach for markov random fields. In: *ECCV*, vol. 7577, pp. 386–399. Springer Berlin Heidelberg (2012)