Fast and Exact Primal-Dual Iterations for Variational Problems in Computer Vision

Jan Lellmann Dirk Breitenreicher Christoph Schnörr
Image and Pattern Analysis Group & HCI, University of Heidelberg
e-mail: {lellmann, breitenreicher, schnoerr}@math.uni-heidelberg.de

Overview

- Optimization for *convex non-smooth* variational problems: relaxed combinatorial problems, lifted functionals, ...
- Based on generic saddle point formulation
- Dual Multiple-Constraint Douglas-Rachford (DMDR) optimization
- Handles difficult dual constraints $\Rightarrow$ tighter relaxation possible
- Globally convergent, order of magnitude faster than FPD or DR
- Bonus: Increased numerical robustness

Saddle-Point Framework

Variational Approach. Define output (image, labels, ...) as minimizer of functional,

$$u^* := \arg\min_{u \in \mathcal{C}} f(u)\,. \tag{1}$$

Many interesting f are convex but nonsmooth (ℓ^1 sparsity, total variation), which complicates optimization. Remedy: *saddle point formulation,*

$$\min_{u \in \mathcal{C}} \max_{v \in \mathcal{D}} g(u,v)\,, \tag{2}$$

$$g(u,v) := \langle u, s\rangle + \langle Lu, v\rangle - \langle b, v\rangle\,, \tag{3}$$

with $s \in \mathbb{R}^n$, $b \in \mathbb{R}^m$, $L \in \mathbb{R}^{m\times n}$, closed convex sets $\mathcal{C} \subseteq \mathbb{R}^n$ and $\mathcal{D} \subseteq \mathbb{R}^m$. Frequently, $\mathcal{D}$ is complicated, but is the intersection of simple sets:

$$\mathcal{D} = \mathcal{D}_1 \cap \ldots \cap \mathcal{D}_r\,. \tag{4}$$

Application: Multiclass Labeling

Problem. For each pixel $x \in \Omega \subseteq \mathbb{R}^d$, find a class label $\ell(x) \in \{1,\ldots,l\}$ according to *local data fidelity* and *regularization term.* Penalize boundary length between classes i and j by $d(i,j)$. *Combinatorial* problem!

Convex Relaxation. Embed labels into $\mathbb{R}^l$, identify the i-th label with unit vector e^i and relax to the unit simplex. Find

$$\min_{u\in\mathcal{C}}\{\langle u,s\rangle + J(u)\}\,,\quad \mathcal{C} := \{u \in \mathrm{BV}(\Omega,\mathbb{R}^l) | u(x) \geqslant 0, \sum_i u_i(x) = 1\}\,, \tag{5}$$

$$J(u) := \sup_{v\in\mathcal{D}} \int_\Omega \langle Du, v\rangle, \quad \mathcal{D} := \{v \in (C_c^\infty)^{d\times l} | v(x) \in \mathcal{D}_{\mathrm{loc}} \forall x \in \Omega\}\,, \tag{6}$$

$$\mathcal{D}_{\mathrm{loc}} := \{v = (v^1,\ldots,v^l) \in \mathbb{R}^{d\times l} | \|v^i - v^j\| \leqslant d(i,j)\ \forall i<j, \sum_k v^k = 0\}\,. \tag{7}$$

$\Rightarrow$ Tight relaxation, but difficult regularizer/dual constraints: Non-smooth, no closed-form expression. But: fits naturally into saddle point framework.

Douglas-Rachford Splitting for Multiple Dual Constraints

Motivation. First-order methods such as FPD need to compute iterative projections on $\mathcal{D}$, which are slow and inexact.

Approach. Reduce to projections on the $\mathcal{D}_i$ by solving the *dual* problem

$$\max_{v \in \mathcal{D}_1\cap,\ldots,\cap\mathcal{D}_r} \min_{u\in\mathcal{C}} g(u,v)\,, \tag{8}$$

introducing auxiliary variables and splitting according to

$$\min_{v_i\in\mathbb{R}^m} \underbrace{\delta_{-L^\top(\frac{1}{r}\sum_i v_i)=z, v_1=\ldots=v_r}}_{f_1} + \underbrace{\textstyle\sum_i \delta_{v_i\in\mathcal{D}_i} + \langle \frac{1}{r}\sum_i v_i, b\rangle + \max_{u\in\mathcal{C}}\langle u, z-s\rangle}_{f_2}\,. \tag{9}$$

Apply Douglas-Rachford iteration, which requires proximal steps on f_1 and f_2, and converges for any step size. But: How to extract a primal solution u?

Algorithm and Properties

Algorithm 1 Dual Multiple-Constraint Douglas-Rachford Optimization for Saddle-Point Problems (DMDR)

```
Choose τ > 0, v̄_i^0 ∈ R^{n×d×l}, z̄^0 ∈ R^{n×d}. Set k ← 0.
while (not converged) do
    v_i^k ← Π_{D_i}(v̄_i^k − (τ/r) b).
    z''^k ← Π_C((1/τ)(z̄^k − s)).
    v'^k ← (rI + LL^⊤)^{−1} (Σ_i (2v_i^k − v̄_i^k) − L(z̄^k − 2τ z''^k)).
    v_1'^k = ... = v_r'^k ← v'^k.
    z'^k ← (−L^⊤) v'^k.
    v̄_i^{k+1} ← v̄_i^k + v_i'^k − v_i^k.
    z̄^{k+1} ← z'^k + τ z''^k.
    k ← k + 1.
end while
```

Proposition 1. Let $\mathcal{D}_1,\ldots,\mathcal{D}_r,\mathcal{C}$ be closed convex sets, $\mathcal{C}$ bounded and $\mathrm{ri}(\mathcal{D}_1)\cap\ldots\cap\mathrm{ri}(\mathcal{D}_r) \neq \emptyset$ and $\mathrm{ri}(\mathcal{C}) \neq \emptyset$. Then Alg. 1 converges in (v_i^k, z''^k).

Proposition 2. Let $(v := v_1 = \ldots = v_r, z'')$ be a fixed point of Alg. 1. Then z'' is a solution of the primal problem (2).

Properties.

- First-order method: suitable for large scale, easy to exploit sparsity.
- Requires only projections on the simple sets $\mathcal{C}$ and $\mathcal{D}_i$.
- Linear equation system can often be solved explicitly using DCT.
- Convergent for any step size τ.
- Handles large class of variational problems.

References

- Chambolle, A., Cremers, D., Pock, T.: A convex approach for computing minimal partitions. Tech. Rep. 649, Ecole Polytechnique CMAP (2008)
- Douglas, J., Rachford, H.H.: On the numerical solution of heat conduction problems in two and three space variables. Trans. of the AMS **82** (1956) 421–439
- Lellmann, J., Schnörr, C.: Continuous multiclass labeling approaches and algorithms. Tech. rep., Univ. of Heidelberg (2010)
- Pock, T., Cremers, D., Bischof, H., Chambolle, A.: An algorithm for minimizing the Mumford-Shah functional. In: Int. Conf. Comp. Vis. (2009)

Experiments

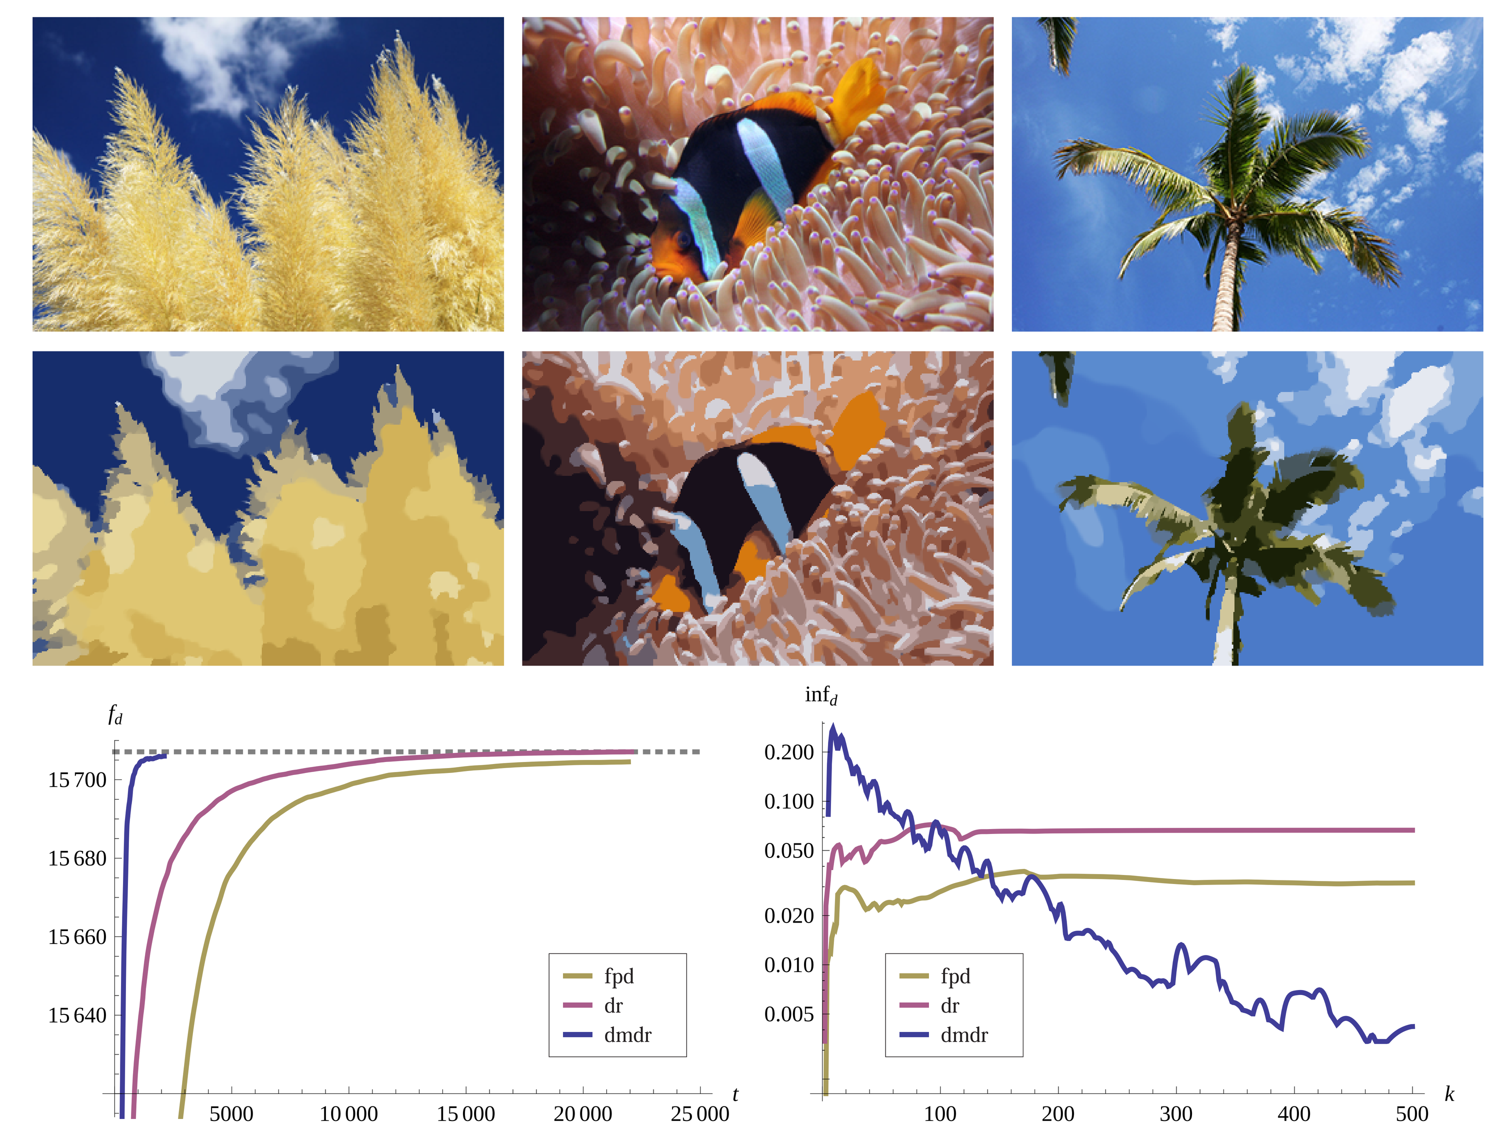

Figure: Runtime performance on labeling problems with 12 classes. **Top row:** Input images (top) and segmentation into 12 classes (bottom). **Left:** Dual objective vs. time for 500 iterations on the "crop" image. The proposed method outperforms DR and FPD by a factor of 10 resp. 17. **Right:** Infeasibility of the dual iterates vs. number of iterations. Due to inexact projections, FPD and DR get stuck and converge to infeasible solutions, while DMDR gradually decreases the infeasibility to zero, increasing numerical robustness.

Experiments

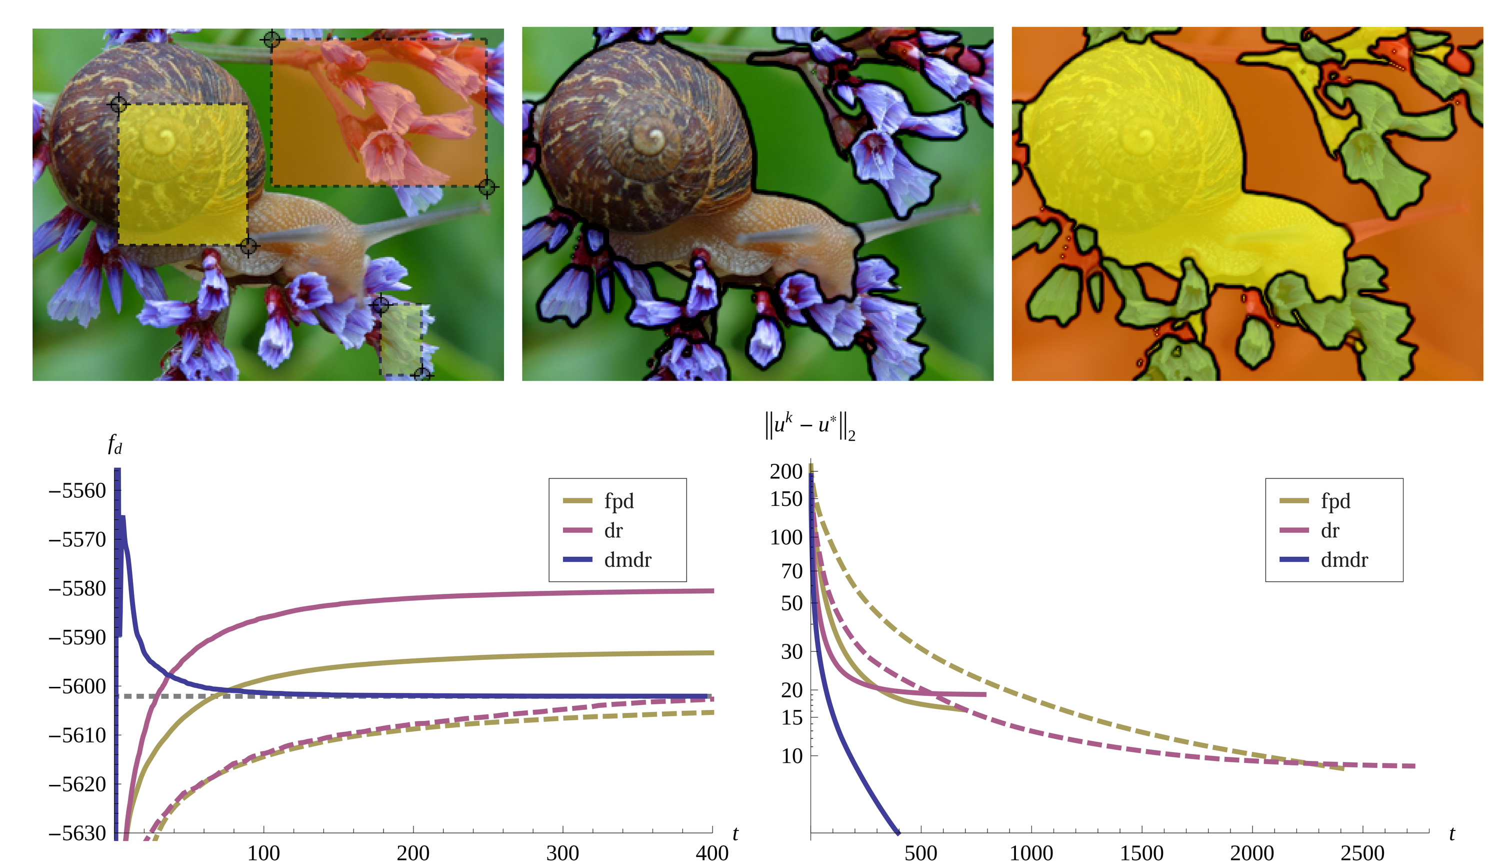

Figure: Application to multiclass segmentation of weakly labeled input image. **Top row:** Segmentation of input image into 3 classes based on seed regions marked by the user. **Bottom row:** Dual objectives (left) and ℓ_2 distance to the reference solution (right) vs. time. With low-accuracy approximate projections, FPD and DR get stuck in an infeasible solution (solid). Increasing the projection accuracy reduces the effect but slows down convergence (dashed). The proposed DMDR method avoids these problems and returns high-quality solutions after only a few iterations.