An Alternating Direction Method for Optical Flow Estimation with l_p Regularization

Naftali Zon
School of Electrical Engineering
Tel Aviv University, Israel

Nahum Kiryati
School of Electrical Engineering
Tel Aviv University, Israel

***Abstract*—We consider the optical flow estimation problem with l_p sub-quadratic regularization, where $0 \leq p \leq 1$. As in other image analysis tasks based on functional minimization, sub-quadratic regularization is expected to admit discontinuities and avoid oversmoothing of the estimated optical flow field. The problem is mathematically challenging, since the regularization term is non-differentiable. It is harder than the l_1 case, that can be addressed via Moreau proximal mapping with a closed form solution. In this paper, we propose a novel approach, based on variable splitting and the Alternating Direction Method of Multipliers (ADMM). We exemplify that our method can outperform optical flow with l_1 regularization, but this is not the essence of this paper. The contribution is in demonstrating that state of the art optimization methods can be harnessed to solve a mathematically-challenging class of important image processing problems, and to highlight crucial numerical aspects that are often obscured in the image processing literature.**

I. Introduction

The field of optical flow estimation experiences constant progress as seen in state of the art results on the Middlebury [1], KITTI 2012 [2], KITTI 2015 [3] and SINTEL [4] optical flow benchmarks. While recent deep-learning based approaches show great promise [5], most optical flow algorithms are descendants of the seminal approaches of Lucas-Kanade [6] or Horn-Schunck [7].

The Horn-Schunck method is the root of optical flow algorithms that are based on functional minimization. Its basis is the brightness-constancy assumption

$$I(\mathbf{x},t) = I(\mathbf{x}+\delta\mathbf{x}, t+\delta t), \tag{1}$$

where $I(\mathbf{x},t)$ is the image brightness at position $\mathbf{x}$ and time t. Optical flow is defined as the apparent motion field between consecutive frames in an image sequence:

$$\mathbf{w} = \begin{bmatrix} u \\ v \end{bmatrix} \equiv \frac{\delta \mathbf{x}}{\delta t} \tag{2}$$

The aperture problem demonstrates that optical-flow estimation is locally under-constrained, calling for a-priori assumptions on the smoothness of the optical flow field $\mathbf{w}$. The functional minimization approach to optical flow estimation expresses the brightness-constancy assumption as a fidelity term, and the a-priori smoothness assumptions as a regularization term, to obtain a functional such as

$$E(u,v) = \frac{\gamma}{2} \left\| I_2(x+u, y+v) - I_1(x,y) \right\|^2 + \left\| \nabla u \right\|^2 + \left\| \nabla v \right\|^2, \tag{3}$$

ideally minimized by the latent optical flow $\mathbf{w}$.

The objective functional (3) has many variations. For example, an additional gradiant-constancy fidelity term [8]

$$\nabla I(\mathbf{x},t) = \nabla I(\mathbf{x}+\delta\mathbf{x}, t+\delta t) \tag{4}$$

or replacement of quadratic regularization by an anisotropic diffusion term [9], [10]. Interestingly, Sun *et al*. [11] discovered that the Brox *et al*. [8] and Horn-Schunck [7] formulations perform surprisingly well when combined with modern optimization techniques and ad-hoc amendments. One key alteration is median filtering of intermediate flow fields during optimization. While this improves the robustness of classical methods, it actually leads to higher energy solutions, meaning that the original objective function is compromized.

II. l_p-norm Regularization

The study of regularization terms has been central to the development of image restoration algorithms. Image restoration with quadratic penalty (l_2-norm) on image derivatives is simple and efficient, but tends to over-smooth the result. Thus, modern edge-preserving image restoration methods generally incorporate sub-quadratic regularization terms. Common approaches incorporate the isotropic and anisotropic (l_1-norm) TV-regularization [12], [13], and l_p-norms with $0 < p < 1$ [14], [15].

The analogy between optical flow field discontinuities and grey-level edges encourages migration from quadratic to sub-quadratic regularization terms in optical flow estimation as well. Along these lines, optical flow estimation algorithms with l_1 regularization have been suggested [21], [22]. Neverthess, optical flow estimation with the potentially superior but mathematically challenging l_p regularization, where $0 \leq p < 1$, has remained an open problem.

In this paper we present a novel optical flow estimation scheme, with l_p-norm regularization, $0 \leq p \leq 1$. Exdending Eq. 3, the generic optical flow functional can be expressed as

$$\begin{aligned} E(u,v) = & \frac{\gamma}{2} \left\| I_2(x+u, y+v) - I_1(x,y) \right\|^2 \\ & + R_u(\nabla u) + R_v(\nabla v) \end{aligned} \tag{5}$$

where R_u, R_v are the regularization functions. For $\|(u,v)\| << 1$ this functional can be expressed as

$$E(u,v) = \frac{\gamma}{2} \left\| I_x u + I_y v + I_t \right\|^2 + R_u(\nabla u) + R_v(\nabla v) \tag{6}$$

978-1-5090-2152-9/16/$31.00

Functional (6) is merely an approximation of functional (5). Minimization of (5) can be carried out using a successive approximation process, in which the optical flow found by minimization of (6) is substituted in (5), the derivatives I_x, I_y and I_t are re-evaluated, the residual flow is recomputed using (6), and so on until convergence. See [16] for the rigorous justification.

The minimization of optical flow objective functionals is usually addressed using the Euler-Lagrange variational approach [8]. This is a difficult mathematical exercise when the functional is not smooth. It can be addressed by smooth approximations of the regularization terms. In this paper, we follow an alternative approach. We discretize the objective functional (6), and derive the optical flow estimation algorithm directly from the discrete version. As will be seen, modern optimization tools facilitate the minimization of the non-smooth regularizer without smooth approximation. We vectorize u, v into $\mathbf{u}, \mathbf{v}$ and set $\mathbf{I_x} = \text{diag}(\text{vec}(I_x))$, $\mathbf{I_y} = \text{diag}(\text{vec}(I_y))$, $\mathbf{I_t} = \text{vec}\{I_t\}$ to be diagonal matrices where the diagonals are the image derivatives I_x, I_y and I_t. We denote by $\mathbf{D_x}$ and $\mathbf{D_y}$ the matrices corresponding to x and y-derivative filters, e.g. $\mathbf{D}_x\mathbf{u} = \text{vec}\left\{u * [\ 1 \quad -1\ ]\right\}$. The continuous objective (6) can now be expressed in column stack discrete form as

$$\begin{aligned} E(\mathbf{u},\mathbf{v}) =& \frac{\gamma}{2}||\mathbf{I_x}\mathbf{u} + \mathbf{I_y}\mathbf{v} + \mathbf{I_t}||^2 \\ &+\Phi(\mathbf{D_x}\mathbf{u}, \mathbf{D_y}\mathbf{u}) + \Phi(\mathbf{D_x}\mathbf{v}, \mathbf{D_y}\mathbf{v}), \end{aligned} \tag{7}$$

where $\Phi(\mathbf{r},\mathbf{s}) = \sum_i (r_i^2 + s_i^2)^{1/2}$ is the *isotropic* l_1 norm regularizer [19]. The *isotropic* l_p regularization norm can be expressed as $\Phi(\mathbf{r},\mathbf{s}) = \sum_i (r_i^2 + s_i^2)^{p/2}$ [23]. Note that *anisotropic* l_p regularization was used in [14], [15].

III. ADMM

In this section we review the *Alternating Direction Method of Multipliers* (ADMM) [17]. ADMM is a fast optimization tool, that achieves its high performance by splitting a complicated minimization problem into simpler sub-problems. These sub-problems usually have a close form solution. ADMM will be shown to be an efficient approach to address the minimization of (7).

We present ADMM via the general unconstrained minimization problem

$$\min_{\mathbf{x}}[f_1(\mathbf{x}) + f_2(\mathbf{D}\mathbf{x})], \tag{8}$$

where $\mathbf{x} \in \mathbb{R}^n$ is a minimization variable and $\mathbf{D} \in \mathbb{R}^{n\times n}$ is a matrix. We introduce a variable $\mathbf{y}$ that splits the objective functional (8) to obtain an equivalent constrained problem

$$\min_{\mathbf{x},\mathbf{y}}[f_1(\mathbf{x}) + f_2(\mathbf{y})], \quad \text{s.t.} \quad \mathbf{D}\mathbf{x} = \mathbf{y}\ . \tag{9}$$

The *Augmented Lagrangian* function for the constrained problem (9) is constructed as

$$\mathcal{L}(\mathbf{x},\mathbf{y},\lambda,\alpha) = f_1(\mathbf{x}) + f_2(\mathbf{y}) + \lambda^T(\mathbf{y} - \mathbf{D}\mathbf{x}) + \frac{\alpha}{2}\left\|\mathbf{D}\mathbf{x} - \mathbf{y}\right\|_2^2 \tag{10}$$

where λ is a vector of Lagrange multipliers and α is a penalty parameter [18]. The *Augmented Lagrangian Method* (ALM) solves (9) by iteratively minimizing the augmented Lagrangian function (10) and updating the multipliers λ. ALM is summarized in Alg. 1 using dual Lagrange multipliers $\mathbf{d}$.

Algorithm 1: Augmented Lagrangian Method (ALM)

Input: $\mathcal{L}$
Output: $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ minimizer of (9)
1 Initialization $\mathbf{d}^{(0)} = \mathbf{0}$
2 **for** *l = 1,2,..., until convergence* **do**
3 $\quad (\mathbf{x}^{(l+1)}, \mathbf{y}^{(l+1)}) =$
$\quad\quad \text{argmin}_{\mathbf{x},\mathbf{y}}[f_1(\mathbf{x}) + f_2(\mathbf{y}) + \frac{\alpha}{2}||\mathbf{D}\mathbf{x} - \mathbf{y} - \mathbf{d}^{(l)}||^2]$
4 $\quad \mathbf{d}^{(l+1)} = \mathbf{d}^{(l)} - \mathbf{D}\mathbf{x}^{(l+1)} + \mathbf{y}^{(l+1)}$
5 **end**
6 Return $(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = (\mathbf{x}^{(l)}, \mathbf{y}^{(l)})$

Simultaneous minimization of $\mathbf{x}$ and $\mathbf{y}$ in line 3 of Alg. 1 is nontrivial since it involves the non-separable quadratic term $\frac{\alpha}{2}||\mathbf{D}\mathbf{x} - \mathbf{y} - \mathbf{d}^{(l)}||^2$ in addition to $f_1(\mathbf{x})$ and $f_2(\mathbf{y})$. A natural way to address the minimization problem in line 3 is by alternate minimization with respect to $\mathbf{x}$ and $\mathbf{y}$, while keeping the other fixed. Experimental evidence in [19] suggests that an efficient algorithm is obtained by running just one minimization step with respect to $\mathbf{x}$ and $\mathbf{y}$. The resulting Alternating Direction Method of Multipliers (ADMM) is summarized in Alg. 2.

Algorithm 2: ADMM

Input: $\mathcal{L}$
Output: $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ minimizer of (9)
Initialization $\mathbf{d}^{(0)} = \mathbf{0}$
for *l = 1,2,..., until convergence* **do**
$\quad \mathbf{x}^{(l+1)} = \text{argmin}_{\mathbf{x}}[f_1(\mathbf{x}) + \frac{\alpha}{2}||\mathbf{D}\mathbf{x} - \mathbf{y}^{(l)} - \mathbf{d}^{(l)}||^2]$
$\quad \mathbf{y}^{(l+1)} = \text{argmin}_{\mathbf{y}}[f_2(\mathbf{y}) + \frac{\alpha}{2}||\mathbf{D}\mathbf{x}^{(l+1)} - \mathbf{y} - \mathbf{d}^{(l)}||^2]$
$\quad \mathbf{d}^{(l+1)} = \mathbf{d}^{(l)} - \mathbf{D}\mathbf{x}^{(l+1)} + \mathbf{y}^{(l+1)}$
end
Return $(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = (\mathbf{x}^{(l)}, \mathbf{y}^{(l)})$

IV. Proposed Method

In this section we address the minimization of functional (7) using the ADMM approach. We include crucial implementation details that are rarely discussed in the image processing literature.

Following variable splitting

$$\begin{aligned} \mathbf{u}_x = \mathbf{D}_x\mathbf{u}, \quad \mathbf{u}_y = \mathbf{D}_y\mathbf{u}, \\ \mathbf{v}_x = \mathbf{D}_x\mathbf{v}, \quad \mathbf{v}_y = \mathbf{D}_y\mathbf{v} \end{aligned} \tag{11}$$

minimizing functional (7) takes the form

$$\begin{aligned} E(\mathbf{u},\mathbf{v},\mathbf{u}_x,\mathbf{u}_y,\mathbf{v}_x,\mathbf{v}_y) =& \frac{\gamma}{2}||\mathbf{I_x}\mathbf{u} + \mathbf{I_y}\mathbf{v} + \mathbf{I_t}||^2 \\ &+\Phi(\mathbf{u}_x,\mathbf{u}_y) + \Phi(\mathbf{v}_x,\mathbf{v}_y) \end{aligned} \tag{12}$$

The Augmented Lagrangian function of problem (12) is

$$\begin{aligned}&\mathcal{L}(\mathbf{u},\mathbf{v},\mathbf{u_x},\mathbf{u_y},\mathbf{v_x},\mathbf{v_y},\mathbf{a_x},\mathbf{a_y},\mathbf{b_x},\mathbf{b_y}) = \\ &= \frac{\gamma}{2}||\mathbf{I_x}\mathbf{u}+\mathbf{I_y}\mathbf{v}+\mathbf{I_t}||^2 \\ &+ \Phi(\mathbf{u_x},\mathbf{u_y}) + \Phi(\mathbf{v_x},\mathbf{v_y}) \\ &+ \frac{\alpha}{2}||\mathbf{D_x}\mathbf{u}-\mathbf{u_x}-\mathbf{a_x}||^2 \\ &+ \frac{\alpha}{2}||\mathbf{D_y}\mathbf{u}-\mathbf{u_y}-\mathbf{a_y}||^2 \\ &+ \frac{\alpha}{2}||\mathbf{D_x}\mathbf{v}-\mathbf{v_x}-\mathbf{b_x}||^2 \\ &+ \frac{\alpha}{2}||\mathbf{D_y}\mathbf{v}-\mathbf{v_y}-\mathbf{b_y}||^2,\end{aligned} \tag{13}$$

where $\mathbf{a}_x, \mathbf{a}_y, \mathbf{b}_x, \mathbf{b}_y$ are the Lagrange multipliers vectors and $\alpha \geq 0$ is the penalty weight parameter. Using ADMM, we minimize with respect to each variable independently, then update the Lagrangian multipliers, as sketched in Alg. 3:

Algorithm 3: Optical flow estimation, l_p reg., sketch

Input: $\mathcal{L}$
Output: $(\hat{\mathbf{u}}, \hat{\mathbf{v}})$ minimizer of (7)
for *k = 1,2,..., until convergence* **do**
  $(\mathbf{u}^{(k+1)}, \mathbf{v}^{(k+1)}) = \text{argmin}_{\mathbf{u},\mathbf{v}}\, \mathcal{L}(...)$
  $(\mathbf{u_x}^{(k+1)}, \mathbf{u_y}^{(k+1)}, \mathbf{v_x}^{(k+1)}, \mathbf{v_y}^{(k+1)}) =$ $\text{argmin}_{\mathbf{u_x},\mathbf{u_y},\mathbf{v_x},\mathbf{v_y}}\, \mathcal{L}(...)$
  $(\mathbf{a_x}^{(k+1)}, \mathbf{a_y}^{(k+1)}, \mathbf{b_x}^{(k+1)}, \mathbf{b_y}^{(k+1)}) =$ `updateMultipliers(...)`
end
Return $\hat{\mathbf{u}} = \mathbf{u}^{(k)}, \hat{\mathbf{v}} = \mathbf{v}^{(k)}$

Minimization of $\mathcal{L}$ with respect to $\mathbf{u}, \mathbf{v}$ yields the following linear equations:

$$\begin{aligned}&\gamma \mathbf{I_x}^T(\mathbf{I_x}\mathbf{u}+\mathbf{I_y}\mathbf{v}+\mathbf{I_t}) \\ &+\alpha(\mathbf{D_x}^T(\mathbf{D_x}\mathbf{u}-\mathbf{u_x}-\mathbf{a_x})+\mathbf{D_y}^T(\mathbf{D_y}\mathbf{u}-\mathbf{u_y}-\mathbf{a_y})) = 0 \\ &\gamma \mathbf{I_y}^T(\mathbf{I_x}\mathbf{u}+\mathbf{I_y}\mathbf{v}+\mathbf{I_t}) \\ &+\alpha(\mathbf{D_x}^T(\mathbf{D_x}\mathbf{v}-\mathbf{v_x}-\mathbf{b_x})+\mathbf{D_y}^T(\mathbf{D_y}\mathbf{v}-\mathbf{v_y}-\mathbf{b_y})) = 0\end{aligned}$$

The resulting sparse linear system of equations can be solved with common numerical methods, such as the *Gauss-Seidel* or *SOR* iterative methods.

The advantage of variable splitting and ADMM is that minimization with respect to $\mathbf{u}_x$ and $\mathbf{u}_y$ (minimization with respect to $\mathbf{v}_x$ and $\mathbf{v}_y$ is similar) yields the Moreau proximal mapping [20] applied to $\mathbf{D}_x\mathbf{u}^{(l+1)} - \mathbf{a}_x^{(l)}$ and $\mathbf{D}_y\mathbf{u}^{(l+1)} - \mathbf{a}_y^{(l)}$. Formally, the Moreau proximal mapping function $\Psi_{\alpha,l_p}(\mathbf{c})$ is defined as

$$\Psi_{\alpha,\phi}(\mathbf{c}) = \underset{\mathbf{q}}{\text{argmin}}\, \frac{\alpha}{2}\left\|\mathbf{q}-\mathbf{c}\right\|^2 + \phi(\mathbf{q}) \tag{14}$$

This function has been extensively studied [20], [24] as it is a key ingredient of ISTA algorithms for non-smooth minimization. Knowing the proximity mapping function, we directly obtain the minimization expression for $\mathbf{u_x}, \mathbf{u_y}$:

$$\begin{aligned}\arg\min_{\mathbf{u_x},\mathbf{u_y}} \mathcal{L}() = \arg\min_{\mathbf{u_x},\mathbf{u_y}} \Bigg[&\sum_i ([\mathbf{u_x}]_i^2 + [\mathbf{u_y}]_i^2)^{p/2} \\ &+\frac{\alpha}{2}||\mathbf{D_x}\mathbf{u}^{(l+1)} - \mathbf{u_x} - \mathbf{a_x}^{(l)}||^2 \\ &+\frac{\alpha}{2}||\mathbf{D_y}\mathbf{u}^{(l+1)} - \mathbf{u_y} - \mathbf{a_y}^{(l)}||^2\Bigg],\end{aligned} \tag{15}$$

where, as discussed following Eq. (7), the l_p regularization function $\Phi(\mathbf{u}_x, \mathbf{u}_y)$ is expressed in discrete form as the sum (over all pixels) $\sum_i([\mathbf{u_x}]_i^2 + [\mathbf{u_y}]_i^2)^{p/2}$. This multidimensional minimization can be reduced to 2D minimization for each pixel:

$$\begin{aligned}([\hat{\mathbf{u}}_\mathbf{x}]_i, [\hat{\mathbf{u}}_\mathbf{y}]_i) = \arg\min_{[\mathbf{u_x}]_i,[\mathbf{u_y}]_i} \Big[&([\mathbf{u_x}]_i^2 + [\mathbf{u_y}]_i^2)^{p/2} \\ &+ \frac{\alpha}{2}([\mathbf{D}_x\mathbf{u}^{(l+1)} - \mathbf{a_x}^{(l)}]_i - [\mathbf{u_x}]_i)^2 \\ &+ \frac{\alpha}{2}([\mathbf{D}_y\mathbf{u}^{(l+1)} - \mathbf{a_y}^{(l)}]_i - [\mathbf{u_y}]_i)^2\Big]\end{aligned} \tag{16}$$

We define the following variables in $\mathbb{R}^2$:

$$\mathbf{q} = \begin{bmatrix}[\mathbf{u_x}]_i \\ [\mathbf{u_y}]_i\end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix}\left[\mathbf{D}_x\mathbf{u}^{(l+1)} - \mathbf{a_x}^{(l)}\right]_i \\ \left[\mathbf{D}_y\mathbf{u}^{(l+1)} - \mathbf{a_y}^{(l)}\right]_i\end{bmatrix} \tag{17}$$

The solution to the 2D minimization problem (16) is equivalent to the following Moreau proximal mapping:

$$\Psi_{\alpha,l_p}(\mathbf{c}) = \underset{\mathbf{q}}{\text{argmin}}\, \frac{\alpha}{2}\left\|\mathbf{q}-\mathbf{c}\right\|_2^2 + \left\|\mathbf{q}\right\|_2^p \tag{18}$$

For the l_1-norm, the corresponding proximity mapping function $\mathbf{\Psi}_{\alpha,l_1}$ has the closed form of soft thresholding

$$\mathbf{\Psi}_{\alpha,l_1}(\mathbf{c}) = \text{sgn}(\frac{\mathbf{c}}{\left\|\mathbf{c}\right\|_2}) \cdot \max\{\left\|\mathbf{c}\right\|_2 - \frac{1}{\alpha}, 0\} \tag{19}$$

For the l_0-norm regularizer, the corresponding proximity mapping function $\mathbf{\Psi}_{\alpha,l_0}$ has the closed form of hard thresholding

$$\mathbf{\Psi}_{\alpha,l_0}(\mathbf{c}) = \frac{\mathbf{c}}{\left\|\mathbf{c}\right\|_2}\mathbf{1}_{\|\mathbf{c}\|_2 \geq \sqrt{2/\alpha}} \tag{20}$$

where $\mathbf{1}$ is an indicator function. For the general l_p-norm regularizer we can calculate $\mathbf{\Psi}_{\alpha,l_p}$ in advance (off-line) for sufficiently many values of $\mathbf{c}$, and store the results in a look-up table (LUT) for use during minimization.

The detailed minimization of (7) using ADMM is summarized in Alg. 4.

We exemplify the operation of the proposed algorithm with l_p regularization, $p = 0.3$. As discussed above, we address the minimization problem (5) by successive approximation, i.e. minimization of (6) and re-initialization of the spatial and temporal derivatives of I. This leads to a nested iterative process, where the ADMM iterations are employed in the internal loop to minimize (6), and I_x, I_y and I_t are updated in the external loop. In practice, we ran the external loop 5 times,

Algorithm 4: Optical flow estimation, l_p reg., detailed

Input: $\mathcal{L}$

Output: $(\hat{\mathbf{u}}, \hat{\mathbf{v}})$ minimizer of of (7)

for *k = 1,2,..., until convergence* **do**

$(\mathbf{u}^{(k+1)}, \mathbf{v}^{(k+1)}) = \text{argmin}_{\mathbf{u},\mathbf{v}} \mathcal{L}(...)$

$\left([\mathbf{u_x}^{(k+1)}]_i, [\mathbf{u_y}^{(k+1)}]_i\right) = \Phi_{\alpha,l_p}([\mathbf{D_x u}^{(k+1)} - \mathbf{a_x}^{(k)}]_i, [\mathbf{D_y u}^{(k+1)} - \mathbf{a_y}^{(k)}]_i)$

$\left([\mathbf{v_x}^{(k+1)}]_i, [\mathbf{v_y}^{(k+1)}]_i\right) = \Phi_{\alpha,l_p}([\mathbf{D_x v}^{(k+1)} - \mathbf{b_x}^{(k)}]_i, [\mathbf{D_y v}^{(k+1)} - \mathbf{b_y}^{(k)}]_i)$

$\mathbf{a_x}^{(k+1)} = \mathbf{a_x}^{(k)} - \mathbf{D_x u}^{(k+1)} + \mathbf{u_x}^{(k+1)}$

$\mathbf{a_y}^{(k+1)} = \mathbf{a_y}^{(k)} - \mathbf{D_y u}^{(k+1)} + \mathbf{u_y}^{(k+1)}$

$\mathbf{b_x}^{(k+1)} = \mathbf{b_x}^{(k)} - \mathbf{D_x v}^{(k+1)} + \mathbf{v_x}^{(k+1)}$

$\mathbf{b_y}^{(k+1)} = \mathbf{b_y}^{(k)} - \mathbf{D_y v}^{(k+1)} + \mathbf{v_y}^{(k+1)}$

Return $\hat{\mathbf{u}} = \mathbf{u}^{(k)}, \hat{\mathbf{v}} = \mathbf{v}^{(k)}$

with 15 ADMM internal-loop iterations per each execution of the external loop.

Table I presents quantitative comparison between optical flow estimation using TV-l_1 regularization [21], [22] and the proposed method on the *RubberWhale* image from the Middlebury dataset [1]. The comparison is in terms of the standard Average Angular Error (AAE) and End Point Error (EPE) metrics. Fig. 4 is visual comparison between the estimated optical flow using the suggested algorithm and the ground-truth. Figs. 1, 2 and 3 respectively show the progress of the objective function, AAE and EPE values along the iterative process.

Method	AAE	EPE
TV-l_1 [21], [22]	6.9	0.22
Proposed l_p method, $p = 0.3$	6.5	0.22

TABLE I
QUANTITATIVE COMPARISON OF OPTICAL FLOW ESTIMATION PERFORMANCE ON THE *RubberWhale* IMAGE FROM THE MIDDLEBURY DATASET [1]. 1ST ROW: TV-l_1 REGULARIZATION [21], [22]. 2ND ROW: PROPOSED METHOD.

V. CONCLUSION

Optical flow estimation by functional minimization with l_p regularization, where $0 \leq p \leq 1$, is a fundamental image processing task. We presented a novel solution to this problem, directly addressing the non-smooth l_p regularization. We are not aware of previous works employing the l_p norm in the context of optical flow. In a wider context, the suggested approach breaks away from approaches that are based on smooth approximations of the l_p norm. Based on direct discretization of the objective functional, and employing variable splitting and ADMM, we presented an algorithm that is simple to implement. It is much easier than solving the problematic Euler-Lagrange equations associated with the functional, see e.g. [8], [7].

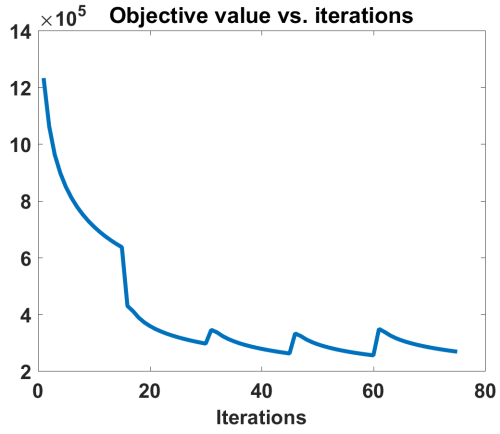

Fig. 1. The objective value exhibits a jumpy decline along the iterative process. The discontinuity following each sequence of 15 ADMM iterations is due to the re-evaluation of I_x, I_y and I_t as the external loop is restarted.

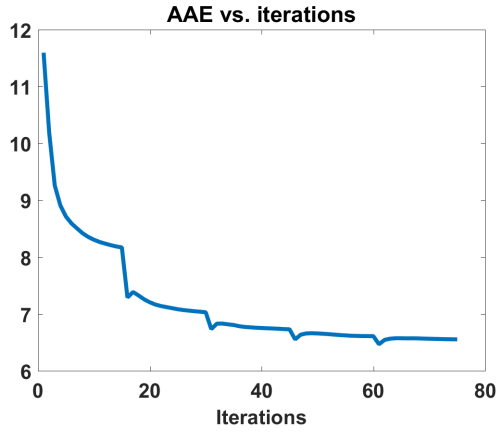

Fig. 2. The declining trend of the AAE with iterations.

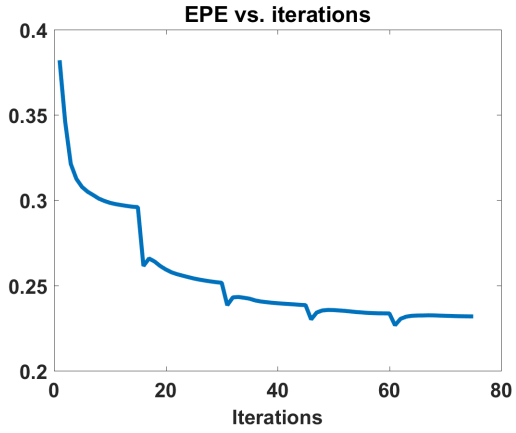

Fig. 3. The progress of the average EPE with iterations.

This work aims to show how state of the art optimization methods can be used to address an important class of image processing problems. Towards this end, it includes algorithmic details that are hard to find in the image processing literature.

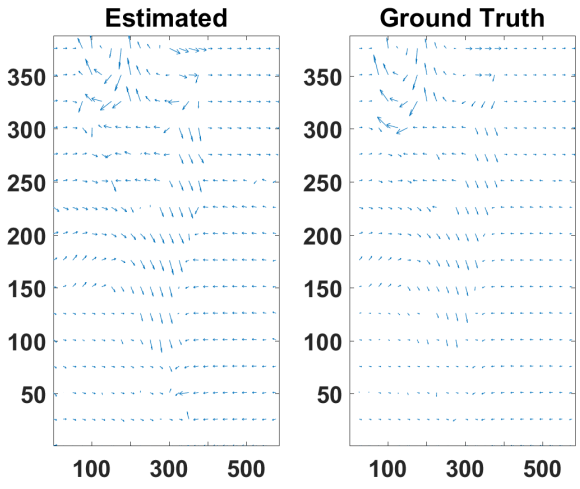

Fig. 4. Comparison of the estimated optical flow to the ground-truth (zoomed-in). Vector coding of vector field ($\mathbf{w}$) (needle map): direction is coded by arrow direction and length by arrow size.

We exemplified the operation of the algorithm, and noted the potential advantage of l_p regularization with respect to l_1 regularizaton.

Obtaining the best possible scores in optical flow benchmarks was not the goal of this research. To approach state-of-the-art performance, one should consider adding a gradient constancy fidelity term, using a coarse-to-fine strategy in the minimization algorithm, and employing various ad-hoc algorithmic and numerical elements as in [11]. We have not incorporated these extensions, as they would have obscured the essentials of this work, and leave them for future research.

REFERENCES

[1] S. Baker, D. Scharstein, J. P. Lewis, S. Roth, M. J. Black, and R. Szeliski, "A database and evaluation methodology for optical flow," *IJCV*, vol. 92, pp. 1–31, 2011.

[2] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2012.

[3] M. Menze, C. Heipke, and A. Geiger, "Joint 3d estimation of vehicles and scene flow," *ISPRS Workshop on Image Sequence Analysis (ISA)*, 2015.

[4] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black, "A naturalistic open source movie for optical flow evaluation," *European Conf. on Computer Vision (ECCV)*, pp. 611–625, 2012.

[5] D. Gadot and L. Wolf, "Patchbatch: a Batch Augmented Loss for Optical Flow," *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[6] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," *Proceedings of the International Joint Conference on Artificial Intelligence*, vol. 2, pp. 674–679, 1981.

[7] B. K. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, pp. 185–203, 1981.

[8] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert, "High accuracy optical flow estimation based on a theory for warping," *Proc. ECCV, Springer LNCS*, vol. 3024, pp. 25–36, 2004.

[9] H. H. Nagel and W. Enkelmann, "An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, pp. 565–593, 1986.

[10] J. Weickert and T. Brox, "Diffusion and regularization of vector- and matrix-valued images," *Inverse Problems, Image Analysis, and Medical Imaging*, pp. 251–268, Dec 2002.

[11] D. Sun, S. Roth, and M. Black, "A quantitative analysis of current practices in optical flow estimation and the principles behind them," *International Journal of Computer Vision*, vol. 106, pp. 115–137, 2014.

[12] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, pp. 259–268, 1992.

[13] Y. Wang, J. Yang, W. Yin, and Y. Zhang, "A new alternating minimization algorithm for total variation image reconstruction," *SIAM Journal on Imaging Sciences*, vol. 1, pp. 248–272, 2008.

[14] A. Levin, R. Fergus, F. Durand, and W. T. Freeman, "Image and depth from a conventional camera with a coded aperture," *ACM Trans. Graph.*, vol. 26, 2007.

[15] D. Krishnan and R. Fergus, "Fast image deconvolution using hyper-laplacian priors," *Neural Information Processing Systems (NIPS)*, 2009.

[16] D. J. Fleet and Y. Weiss, "Optical flow estimation," 2005.

[17] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, pp. 293–318, 1992.

[18] J. Nocedal and S. J. Wright, "Numerical optimization," 2*nd ed. Springer*, 2006.

[19] T. Goldstein and S. Osher, "The split bregman method for L1-regularized problems," *SIAM Journal on Imaging Sciences*, vol. 2, pp. 323–343, 2009.

[20] P. L. Combettes and V. R. Wajs, "Signal recovery by proximal forward-backward splitting," *Multiscale Modeling & Simulation*, vol. 4, pp. 1168–1200, 2005.

[21] C. Zach, T. Pock, and H. Bischof, "A duality based approach for realtime TV-L_1 optical flow," *Proc. DAGM Conference on Pattern Recognition, Springer LNCS*, vol. 4713, pp. 214–223, 2007.

[22] J. Sánchez, E. Meinhardt-Llopis, and G. Facciolo, "TV-L_1 optical flow estimation," *Image Processing On Line*, vol. 3, pp. 137–150, 2013.

[23] J. Kotera, F. Sroubek, and P. Milanfar, "Blind deconvolution using alternating maximum a posteriori estimation with heavy-tailed priors," *Computer Analysis of Images and Patterns (CAIP), LNCS*, vol. 8048, 2013.

[24] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, pp. 127–239, 2014.