
A Quantitative Analysis of Current Practices in Optical Flow Estimation and the Principles behind Them

Deqing Sun · Stefan Roth · Michael J. Black

the date of receipt and acceptance should be inserted later

Abstract The accuracy of optical flow estimation algorithms has been improving steadily as evidenced by results on the Middlebury optical flow benchmark. The typical formulation, however, has changed little since the work of Horn and Schunck. We attempt to uncover what has made recent advances possible through a thorough analysis of how the objective function, the optimization method, and modern implementation practices influence accuracy. We discover that “classical” flow formulations perform surprisingly well when combined with modern optimization and implementation techniques. One key implementation detail is the 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 these methods are not optimizing the original objective function. To understand the principles behind this phenomenon, we derive a new objective function that formalizes the median filtering heuristic. This objective function includes a non-local smoothness term that robustly integrates flow estimates over large spatial neighborhoods. By modifying this new term to include information about flow and image boundaries we develop a method that can better preserve motion details.

Deqing Sun
Department of Computer Science, Harvard University, Cambridge, MA, USA
e-mail: dqsun@seas.harvard.edu
The work for this paper was performed while DS was at Brown University

Stefan Roth
Department of Computer Science, TU Darmstadt, Darmstadt, Germany
e-mail: sroth@cs.tu-darmstadt.de

Michael J. Black
Max Planck Institute for Intelligent Systems, Tübingen, Germany
e-mail: black@is.mpg.de

To take advantage of the trend towards video in wide-screen format, we further introduce an asymmetric pyramid down-sampling scheme that enables the estimation of longer range horizontal motions. The methods are evaluated on Middlebury, MPI Sintel, and KITTI datasets using the same parameter settings.

Keywords Optical flow estimation · Practices · Median filtering · Non-local term · Motion boundary

1 Introduction

The field of optical flow estimation is making steady progress as evidenced by the increasing accuracy of current methods on the Middlebury optical flow benchmark [11]. After nearly 30 years of research, these methods have obtained an impressive level of reliability and accuracy [66,67,69,73,77]. *But what has led to this progress?* The majority of today’s methods strongly resemble the original formulation of Horn and Schunck (**HS**) [31]. They combine a data term that assumes constancy of some image property with a spatial term that models how the flow is expected to vary across the image. An objective function combining these two terms is then optimized. Given that this basic structure is unchanged since **HS**, what has enabled the performance gains of modern approaches?

The paper has three parts. In the first, we perform a study of recent optical flow methods and models. The most accurate methods on the Middlebury flow dataset make different choices about how to model the objective function, how to approximate this model to make it computationally tractable, and how to optimize it. Since most published methods change *all* of these properties at once, it can be difficult to know which choices are most important. To address this, we define a baseline algorithm that is “classical”, in that it is a direct descendant of the original **HS** formulation,

and then systematically vary the model and method using different techniques from the art. The results are surprising. We find that only a small number of key choices produce statistically significant improvements and that they can be combined into a very simple method that achieves reasonable accuracy. More importantly, our analysis reveals what makes current flow methods work so well.

Part two examines the *principles* behind this success. We find that one algorithmic choice produces the most significant improvements: applying a median filter to intermediate flow values during incremental estimation and warping [66, 67]. While this heuristic improves the accuracy of the recovered flow fields, it actually *increases* the energy of the objective function. This suggests that what is being optimized is actually a new and different objective. Using observations about median filtering and L1 energy minimization from Li and Osher [40], we formulate a new *non-local term* that is added to the original, classical objective. This new term goes beyond standard local (pairwise) smoothness to robustly integrate information over large spatial neighborhoods. We show that minimizing this new energy approximates the original optimization with the heuristic median filtering step. Note, however, that the new objective falls outside our definition of classical methods.

Once the median filtering heuristic is formulated as a non-local term in the objective, we immediately recognize how to modify and improve it. In part three we show how information about image structure and flow boundaries can be incorporated into a weighted version of the non-local term to prevent over-smoothing across boundaries. By incorporating structure from the image, this weighted version does not suffer from some of the errors produced by median filtering and better preserves motion boundaries. Figure 1 illustrates optical flow estimates for a range of methods from a “basic” **HS** method to our proposed **Classic+NL** method.

Finally we observe that the classical methods all go beyond the original **HS** algorithm by using a spatial pyramid to cope with large motions. The classical pyramid downsamples the image equally in both the horizontal and vertical direction, typically until some minimum image dimension is reached. With today’s wide-aspect ratio video, we point out that an asymmetric approach can be employed resulting in a pyramid that downsamples more in the horizontal direction than in the vertical one. This effectively allows the estimation of larger horizontal motions. This simple change results in significant improvements on the wide-aspect-ratio video in the KITTI [26] and MPI Sintel [23] datasets.

At the time of writing our previous conference paper [57] (March 2010), the resulting approach was ranked 1st in both angular and end-point errors in the Middlebury evaluation. At the writing of this paper (Sep. 2012), the method, **Classic+NL**, ranks 13th in both AAE and EPE. Several recent and high-ranking methods directly build on **Classic+NL**,

such as layered models [59, 60, 61], methods with more advanced motion prior models [25, 34], efficient optimization schemes for the non-local term [36], and better initialization to deal with large displacement optical flow [24].

Compared to the conference version [57], this paper includes many more detailed results and analyses. In addition to an expanded literature review we compare our proposed method to the closely related non-local total variation method [68]. We discuss the limitations of our method in dealing with occlusions and fast moving objects. We report results on the MIT HAMA data set [42] and find that the results are consistent with those on Middlebury. We also test our methods on the MPI Sintel [23] and KITTI [26] datasets, which offer greater challenges. Using the same parameters, tuned on the Middlebury training set, our method performs well on these new datasets, particularly using an asymmetric pyramid.

In summary, the contributions of this paper are to (1) analyze current flow models and methods to understand which design choices matter; (2) formulate and compare several classical objectives descended from **HS** using modern methods; (3) formalize one of the key heuristics and derive a new objective function that includes a non-local spatial smoothness term; (4) modify this new objective to produce a state-of-the-art method; (5) extend spatial pyramids to exploit the extra width of high-definition and letterbox videos. In doing so, we provide a “recipe” for others studying optical flow that can guide their design choices. Finally, to enable comparison and further innovation, we provide a public MATLAB implementation [1].

2 Previous Work

It is important to separately analyze the contributions of the objective function that defines the problem (*the model*) and the optimization algorithm and implementation used to minimize it (*the method*). The **HS** formulation, for example, has long been thought to be highly inaccurate. Barron *et al.* [12] reported an average angular error (AAE) of ~30 degrees on the “Yosemite” sequence. This confounds the objective function with the particular optimization method proposed by Horn and Schunck. Horn and Schunck noted that the correct way to optimize their objective is by solving a system of linear equations as is common today. This was impractical on the computers of the day, hence they used a heuristic method. In fact, Barron *et al.* note that the original **HS** derivatives were implemented crudely and report a modified version of **HS** with AAE around 11 degrees. When optimized with today’s methods, the **HS** objective achieves surprisingly competitive results [26] despite the expected over-smoothing and sensitivity to outliers. The reported accuracy of a method is jointly determined by the objective function, the optimization techniques, the implementation details, and

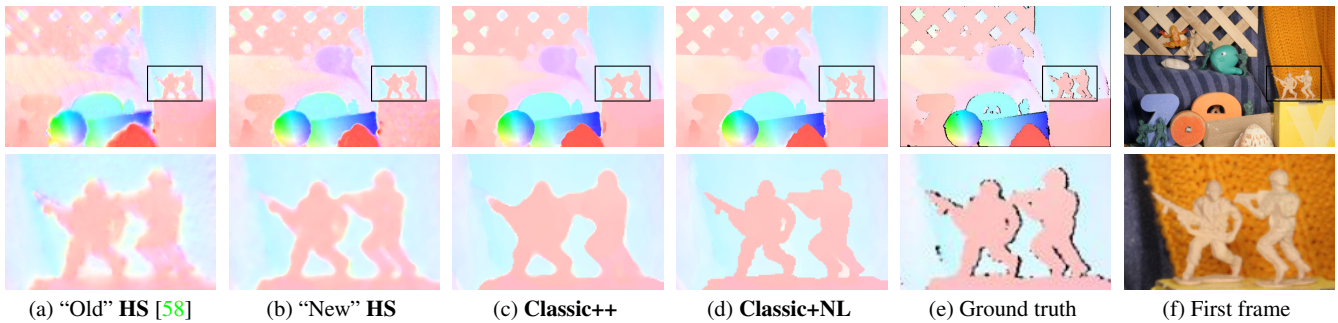


Fig. 1 Estimated optical flow on the Middlebury test “Army” sequence. Left to right: (a) an old implementation of the Horn & Schunck (**HS**) method [58], (b) a new implementation with current practices, (c) a modern implementation of a robust version, (d) an improved model that uses a non-local spatial term to robustly integrate information over a large spatial neighborhood, (e) ground truth from the Middlebury website (downsampled and JPEG compressed; original ground truth is withheld), and (f) the first frame. Color coding as in [11], shown in Fig. 4 (c). Average end-point error (EPE): (a) 0.22, (b) 0.12, (c) 0.09, and (d) 0.08.

the parameter tuning/learning (cf. [45, 62]). We review related research in the context of the first three aspects below.

Models: The global formulation of optical flow introduced by Horn and Schunck [31] relies on both brightness constancy and spatial smoothness assumptions, but suffers from the fact that their quadratic formulation is not robust to outliers. Shulman and Herve [55] use an L1 penalty instead to preserve flow discontinuities. Black and Anandan [15] introduce a robust framework to deal with outliers in both the data and the spatial terms. Subsequently, many different robust functions have been explored [17, 38, 58] and it remains unclear which is best. We refer to all these spatially-discrete formulations derived from **HS** as “classical.” We systematically explore variations in the formulation and optimization of these approaches. The surprise is that the classical model, appropriately implemented, remains fairly competitive.

There are many formulations beyond the classical ones that we do not consider here. Significant ones use oriented smoothness [46, 58, 66, 76, 77], rigidity constraints [65, 66], an over-parameterized smoothness term [48], or image segmentation [14, 37, 72, 78]. While they deserve similar careful consideration, we expect many of our conclusions to carry forward. Note that one can select among a set of models or methods for a given sequence [44], instead of finding a “best” model for all the sequences.

Methods: Many of the implementation details that are thought to be important date back to the early days of optical flow. Current best practices include coarse-to-fine estimation to deal with large motions [13, 17], texture decomposition [65, 67] or high-order filter constancy [8, 17, 29, 39, 77] to reduce the influence of lighting changes, incremental warping [13], warping with bicubic interpolation [38, 67], temporal averaging of image derivatives [30, 67], graduated non-convexity [16] to minimize non-convex energies [15, 58], and median filtering after each incremental estimation step to remove outliers [67].

This median filtering heuristic is of particular interest as it makes non-robust methods more robust and improves the accuracy of all methods we tested. The effect on the objective function and the underlying reason for its success have not previously been analyzed. Least median squares estimation can be used to robustly reject outliers in flow estimation [10], but previous work has focused on the data term.

Related to median filtering, and our new non-local term, is the use of bilateral filtering to prevent smoothing across motion boundaries [71]. This approach separates a variational method into two filtering update stages, and replaces the original anisotropic diffusion process with multi-cue driven bilateral filtering. As with median filtering, the bilateral filtering step changes the original energy function.

Models that are formulated with an L1 robust penalty are often coupled with specialized total variation (TV) optimization methods [75]. Here we focus on generic optimization methods that can apply to most models and find that the estimated flow fields are as accurate as the reported results for specialized methods.

Despite recent algorithmic advances, there is a lack of publicly available, easy to use, and accurate flow estimation software. The GPU4Vision project [2] has made a substantial effort to change this and provides executable files for several accurate methods [65, 66, 67, 69]. The dependence on the GPU and the lack of source code are limitations. Since the publication of our conference paper, our public MATLAB code has been used by both researchers to develop new optical flow algorithms [7, 24, 25, 34, 36] and practitioners to use optical flow for different applications [33, 41, 49]. Currently other available optical-flow software includes [3, 4, 5].

3 Classical Models

As is common to “classical” methods we only address the two-frame optical flow estimation problem. We write the

classical optical flow objective function in its spatially discrete form as

$$E(\mathbf{u}, \mathbf{v}) = \sum_{i,j} \left\{ \rho_D(I_1(i, j) - I_2(i + u_{i,j}, j + v_{i,j})) \right. \quad (1) \\ \left. + \lambda [\rho_S(u_{i,j} - u_{i+1,j}) + \rho_S(u_{i,j} - u_{i,j+1}) \right. \\ \left. + \rho_S(v_{i,j} - v_{i+1,j}) + \rho_S(v_{i,j} - v_{i,j+1})] \right\},$$

where \mathbf{u} and \mathbf{v} are the horizontal and vertical components of the optical flow field to be estimated from images I_1 and I_2 , i, j indexes a particular image pixel location, $u_{i,j}$ and $v_{i,j}$ are elements of \mathbf{u} and \mathbf{v} respectively, λ is a regularization parameter, and ρ_D and ρ_S are the data and spatial penalty functions. We consider three different penalty functions: (1) the quadratic **HS** penalty $\rho(x) = x^2$; (2) the Charbonnier penalty $\rho(x) = \sqrt{x^2 + \epsilon^2}$ [20], a differentiable variant of the absolute value, the most robust convex function; and (3) the Lorentzian $\rho(x) = \log(1 + \frac{x^2}{2\sigma^2})$, which is a non-convex robust penalty used in [15]. We refer to the robust formulation with the Lorentzian penalty as **BA** (short for Black and Anandan). Note that this classical model is related to a standard pairwise Markov random field (MRF) based on a 4-neighborhood [27].

In the remainder of this section we define a baseline method using several techniques from the literature. This is not the “best” method, but includes modern techniques and will be used for comparison. We only briefly describe the main choices, which are explored in more detail in the following section and the cited references.

Quantitative results are presented throughout the remainder of the text. In all cases we report the average end-point error (EPE) on the Middlebury training and test sets, depending on the experiment. Given the extensive nature of the evaluation, only average results are presented in the main body, while the details for each individual sequence are provided at the end of the paper.

3.1 Baseline methods

To gain robustness against lighting changes, we follow [67] and apply the Rudin-Osher-Fatemi (ROF) structure texture decomposition method [52] to pre-process the input sequences and linearly combine the texture and structure components (in the proportion 20:1). The parameters are set according to [67].

Optimization is performed using a standard incremental multi-resolution technique (e. g., [15, 17]) to estimate flow fields with large displacements. The optical flow estimated at a coarse level is used to warp the second image toward the first at the next finer level, and a flow increment is calculated between the first image and the warped second image. The standard deviation of the Gaussian anti-aliasing filter is set to be $\frac{1}{\sqrt{2d}}$, where d denotes the downsampling factor. Each

level is recursively downsampled from its nearest lower level. In building the pyramid, the downsampling factor is not critical as pointed out in the next section; here we use the settings in [58], which uses a factor of 0.8 in the final stages of the optimization. For the basic pyramid scheme, we adaptively determine the number of pyramid levels so that the top level has a width or height of around 20 to 30 pixels. At each pyramid level, we perform 10 warping steps to compute the flow increment.

At each warping step, we linearize the data term once, which involves computing terms of the type $\frac{\partial}{\partial x} I_2(i + u_{i,j}^k, j + v_{i,j}^k)$, where $\partial/\partial x$ denotes the partial derivative in the horizontal direction, u^k and v^k denote the current flow estimate at iteration k . As suggested in [67], we compute the derivatives of the second image using the 5-point derivative filter $\frac{1}{12}[-1 \ 8 \ 0 \ -8 \ 1]$, and warp the second image and its derivatives toward the first using the current flow estimate by bicubic interpolation. We then compute the spatial derivatives of the first image, compute the average of these and the corresponding warped derivatives of the second image (cf. [9, 30]), and use these in place of $\frac{\partial I_2}{\partial x}$. For pixels moving out of the image boundaries, we set both their corresponding temporal and spatial derivatives to zero. After each warping step, the flow update is computed, and then we apply a 5×5 median filter to the newly computed flow field to remove outliers [67].

For the Charbonnier (**Classic-C**) and Lorentzian (**Classic-L**) penalty function, we use a graduated non-convexity (GNC) scheme [16] as described in [58]. First, we replace the robust penalty functions by quadratic penalty functions and obtain a quadratic formulation of the objective function, $E_Q(\mathbf{u}, \mathbf{v})$. Then we linearly combine the quadratic penalty function with the desired robust penalty function and gradually change the weighting of the two terms to reach the desired robust penalty function. In practice, we use a three-stage GNC scheme, with the objective functions for the first, second, and third stages being $E_Q(\mathbf{u}, \mathbf{v})$, $\frac{1}{2}(E_Q(\mathbf{u}, \mathbf{v}) + E(\mathbf{u}, \mathbf{v}))$, and $E(\mathbf{u}, \mathbf{v})$ respectively. The output of a previous stage serves as the initialization to the next stage. The standard deviations of the corresponding quadratic penalty function are set to be 1 for the Charbonnier penalty and, for the Lorentzian, are taken to be the same as the σ value used in the Lorentzian function. The same regularization weight λ is used for both the quadratic and the robust objective functions.

3.2 Baseline results

The regularization parameter λ is selected among a set of candidate values to achieve the best average end-point error (EPE) on the Middlebury training set. For the Charbonnier penalty function, the candidate set is [1, 3, 5, 8, 10] and 5 is optimal. The Charbonnier penalty uses $\epsilon = 0.001$ for both the data and the spatial term in Eq. (1). The Lorentzian uses

	Avg. Rank	Avg. EPE
Classic-C	34.8	0.408
HS	49.0	0.501
Classic-L	42.7	0.530
Classic-C-brightness	N/A	0.726
HS-brightness	N/A	0.759
Classic-L-brightness	N/A	0.603
HS [58]	66.2	0.872
BA (Classic-L) [58]	59.6	0.746
Adaptive [66]	28.5	0.401
Complementary OF [77]	31.6	0.485

Table 1 Models. Average rank and end-point error (EPE) on the Middlebury *test set* using different penalty functions. Two state-of-the-art methods in Dec. 2010 are included for comparison. The ranking information was obtained at the writing of the paper (Sep. 2012). Please refer Table 15 for the EPE results on each sequence.

$\sigma = 1.5$ for the data term, $\sigma = 0.03$ for the spatial term, and $\lambda = 0.06$. These parameters are fixed throughout the experiments, except where mentioned.

Table 1 summarizes the EPE results of the basic model with three different penalty functions on the Middlebury test set, along with the two top performers at the time of performing the evaluation (considering only published papers when the evaluation table was generated). The classic formulations with two non-quadratic penalty functions (**Classic-C**) and (**Classic-L**) achieve competitive results despite their simplicity. The baseline optimization of **HS** and **BA (Classic-L)** results in significantly better accuracy than previously reported for these models [58]. Note that the analysis also holds for the training set (Table 2).

Because **Classic-C** performs quite well despite its simplicity, we set it as the baseline below. Note that our baseline implementation of **HS** has a lower average EPE than many more sophisticated methods. The **HS** implementation here incorporates many algorithmic and implementation details not present in the original **HS** method; the core idea of quadratic data and spatial terms however remains the same. In our naming convention, one can think of the **HS** method here as **Classic-Q**, meaning that it is the same as the **Classic-C** method except that the data and spatial penalty terms are quadratic.

4 Practices Explored

We now systematically vary the baseline approach by incorporating different ideas that have appeared in the literature, with the goal of illuminating which of these ideas are significant. This analysis is performed on the Middlebury training set by changing only *one property at a time*. Statistical significance is determined using a Wilcoxon signed rank test [70] between each modified method and the baseline **Classic-C** method; a p value less than 0.05 indicates a significant difference. Each section below presents detailed comparisons of all these methods and then summarizes the

	Avg. EPE	significance	p -value
Classic-C	0.298	—	—
HS	0.384	1	0.0078
Classic-L	0.319	1	0.0078
Classic-C-brightness	0.288	0	0.9453
HS-brightness	0.387	1	0.0078
Classic-L-brightness	0.325	0	0.2969
Gradient	0.305	0	0.4609
Gaussian + Dx + Dy	0.290	0	0.6406
Sobel edge magnitude [63]	0.417	1	0.0156
Laplacian [39]	0.430	1	0.0078
Laplacian 1:1	0.301	0	0.6641
Gaussian pre-filtering ($\sigma = 0.5$)	0.281	0	0.5469
Texture4:1	0.286	0	0.5312
Unnormalized texture	0.298	0	0.3750

Table 2 Pre-Processing. Average end-point error (EPE) on the Middlebury *training set* for the baseline method (**Classic-C**) using different image pre-processing techniques. Significance is always with respect to **Classic-C**. Please refer to Tables 16 and 17 for the detailed results on each training sequence.

results in a simple “take away message” about what we think are the “best practices” based on the data.

4.1 Image Pre-Processing

While it is common to talk about the brightness constancy assumption as a core feature of most optical flow algorithms, in practice many other constancy assumptions have been used. It is common, for example, to pre-filter the images in a variety of ways ranging from simple smoothing to edge detection. For each method, we optimize the regularization parameter λ for the training sequences. The results are summarized in Table 2, with details of the methods applied to individual training sequences given in Tables 16 and 17. The baseline uses a non-linear pre-filtering of the images (ROF) to reduce the influence of illumination changes between frames [67]. Table 2 shows the effect of using no pre-processing, resulting in the standard brightness constancy model (***-brightness**). **Classic-C-brightness** actually achieves lower EPE on the training set than does **Classic-C** but significantly higher error on the test set (Table 1). This disparity suggests overfitting to the training data and leaves open the question as to whether the standard brightness constancy assumption, formulated robustly, may still compete with various types of filter/structure constancy given appropriate training data.

Simpler alternatives, such as filter response (or high-order) constancy [17, 19, 58] can serve the same purpose as ROF texture decomposition. A variety of pre-filters have been used in the literature, including derivative filters, Laplacians [22, 39], and Gaussians. Edges have also been emphasized using the Sobel edge magnitude [63].

Gradient only imposes constancy of the gradient vector at each pixel as proposed in [17]; i. e., it robustly penalizes Euclidean distance between image gradients. We use central difference filters ($Dx = [-0.5 \ 0 \ 0.5]$ and $Dy = Dx^T$).

Gaussian+Dx+Dy assumes separate brightness, horizontal derivative, and vertical derivative constancy. A weighted combination of robust functions applied to each term is used as in [58]. Neither of these methods differ significantly from the baseline texture decomposition (**Classic-C**). Two methods are significantly worse: the **Sobel edge magnitude** [63] and **Laplacian** pre-filtering (5×5) as used in [39]. **Sobel edge magnitude** appears to not work well on some of the sequences, particularly the synthetic ones, and may not be suitable for a general flow estimation method. **Laplacian** pre-filtering (5×5) as used in [39] produces good results on “RubberWhale”, but poor ones on the synthetic sequences. Note that the parameters for the FusionFlow method [39] were mainly tuned using the “RubberWhale” sequence. The evaluation results suggest room for improving the FusionFlow method by a better pre-processing technique. **Gaussian pre-filtering** ($\sigma = 0.5$) performed well on the synthetic sequences, but poorly on real ones. Finally, the texture-structure blending ratio is 20:1 in [67] but 4:1 in [69]. We find that (**Texture4:1**) performs better (but not significantly) on the synthetic sequences with a little degradation on the real ones. By default, the blended result from texture decomposition is normalized to $[-1, 1]$ in [67] and $[0, 255]$ in our experiment. Not doing this normalization (**Unnormalized texture**) has little effect.

For the Laplacian pre-filtering, we find combining the filtered image with the original image, in the proportion 1:1, improves accuracy significantly (**Laplacian1:1**). Similar to the ROF texture decomposition, such an approach boosts the high frequency while suppressing the low frequency components that contain the lighting change.

Good Practices: Some form of image filtering is useful but simple derivative constancy is nearly as good as the more sophisticated texture decomposition method.

4.2 Coarse-to-Fine Estimation and Graduated Non-Convexity (GNC)

We vary the number of warping steps per pyramid level and find that **3 warping steps** gives similar results as using the baseline 10 (Table 3), except on “Urban3”, which is dominated by large motion and occlusions (see Table 18 for sequence-specific results). For the coarse-to-fine pyramid, [58] uses a downsampling factor of 0.8 during non-convex optimization. A traditional downsampling factor of 0.5 (**Down-0.5**), however, has nearly identical performance. Note that a larger factor means that the pyramid levels are more similar in size and, for a pyramid with top bottom levels of the same size, results in more pyramid levels.

Previously, Brox *et al.* [17] have reported that a downsampling factor of 0.95 produces much better results than 0.5. Note that for each iterative warping estimation step,

	Avg. EPE	significance	p-value
Classic-C	0.298	—	—
3 warping steps	0.304	0	0.9688
Down-0.5	0.298	0	1.0000
w/o GNC	0.354	0	0.1094
Bilinear	0.302	0	0.1016
w/o TAVG	0.306	0	0.1562
Central derivative filter	0.300	0	0.7266
7-point derivative filter [20]	0.302	0	0.3125
Deriv-warp	0.297	0	0.9531
Bicubic-II	0.290	1	0.0391
Deriv-warp-II	0.287	1	0.0156
Warp-deriv-II	0.288	1	0.0391
C-L ($\lambda = 0.6$)	0.303	0	0.1562
L-C ($\lambda = 2$)	0.306	0	0.1562
GC-0.45 ($\lambda = 3$)	0.292	1	0.0156
GC-0.25 ($\lambda = 0.7$)	0.298	0	1.0000
MF 3×3	0.305	0	0.1016
MF 7×7	0.305	0	0.5625
$2 \times$ MF	0.300	0	1.0000
$5 \times$ MF	0.305	0	0.6875
w/o MF	0.352	1	0.0078
Classic++	0.285	1	0.0078

Table 3 Model and Methods. Average end-point error (EPE) on the Middlebury *training set* for the baseline method (**Classic-C**) using different algorithm and modeling choices. Please refer to Table 18 for the detailed results on each sequence.

Brox *et al.* use successive over-relaxation (SOR) to iteratively solve their linear system of equations and stop the iteration before convergence. With a downsampling factor of 0.95, they effectively increase the number of iterative warping steps performed by the algorithm, and this likely helps the overall algorithm converge. For our implementation, we solve the linear system of equations using the MATLAB built-in backslash function and obtain converged results for each iterative warping estimation step. Under such a setting, we find that the downsampling factor has little influence on the performance.

Removing the GNC procedure for the Charbonnier penalty function (**w/o GNC**) results in higher EPE on most sequences and higher energy on all sequences (Table 5). This suggests that the GNC method is helpful even for the convex Charbonnier penalty function due to the nonlinearity of the data term.

Good Practices: The downsampling factor does not matter when using a convex penalty; a standard factor of 0.5 is fine. Some form of GNC is useful even for a convex robust penalty like Charbonnier because of the nonlinear data term.

4.3 Interpolation Method and Derivatives

We find that the baseline bicubic interpolation is more accurate than bilinear (Table 3, **Bilinear**), as already reported in previous work [67]. Removing temporal averaging of the gradients (**w/o TAVG**), using a **Central difference filter** $[-1 \ 0 \ 1]/2$, or using a **7-point derivative filter** $[-1 \ 9 \ -45 \ 0 \ 45 \ -9 \ 1]/60$ [20] all reduce accuracy compared to the baseline, but not significantly.




Fig. 2 Different penalty functions for the spatial terms: Charbonnier ($\epsilon = 0.001$), generalized Charbonnier ($a = 0.45$ and $a = 0.25$), and Lorentzian ($\sigma = 0.03$).

The baseline method computes the image derivative by first computing the derivative of the second image, warping the intermediate result toward the first image, and then averaging the warped result with the spatial derivative of the first image. Another approach is to first warp the second image toward the first image, compute the derivatives of the warped image, and then perform the temporal averaging with the spatial derivatives of the first image [20]. We find the second approach produces similar results (**Deriv-warp**). However, the derivatives computed in either way are inconsistent with those implicitly interpolated by the bicubic interpolation. Bicubic interpolation interpolates not only the image but also the derivatives [50]. Because the MATLAB built-in function `interp2` is based on cubic convolution [35] and does not provide the derivatives used in interpolation, we use the spline-based implementation in [50]. With the new implementation (**Bicubic-II**), the three different ways to compute the derivatives give very similar EPE results, all better than the MATLAB built-in function. However, the one with consistent derivatives (**Bicubic-II**) gives the lowest energy solution, as shown in Table 4.

Good Practices: Use spline-based bicubic interpolation with a 5-point filter. Compute the derivatives during the interpolation to obtain the lowest energy solutions. Temporal averaging of the derivatives is probably worthwhile for a small computational expense.

4.4 Penalty Functions

We find that the convex Charbonnier penalty performs better than the more robust, non-convex Lorentzian on both the training and test sets. We test using the Charbonnier for the data term and Lorentzian for the spatial term (**C-L**) and vice versa (**L-C**). The two approaches perform better than using the Lorentzian for both terms but worse than using the Charbonnier for both terms.

One reason might be that non-convex functions are more difficult to optimize, causing the optimization scheme to find

a poor local optimum. Another reason might be that the MAP estimator actually favors the “wrong” penalty functions [47, 54].

We investigate a generalized Charbonnier penalty function $\rho(x) = (x^2 + \epsilon^2)^a$ that is equal to the Charbonnier penalty when $a = 0.5$, and non-convex when $a < 0.5$ (see Fig. 2). We optimize the regularization parameter λ again. We find a slightly non-convex penalty with $a = 0.45$ (**GC-0.45**) performs consistently better than the Charbonnier penalty, whereas more non-convex penalties (**GC-0.25** with $a = 0.25$) show no improvement.

Good Practices: The less-robust Charbonnier is preferable to the highly non-convex Lorentzian and a slightly non-convex penalty function (**GC-0.45**) is better still.

4.5 Median Filtering

Figure 3 illustrates the median filtering step within the coarse-to-fine incremental estimation process. The baseline 5×5 median filter (**MF** 5×5) is better than both **MF** 3×3 [67] and **MF** 7×7 , but the difference is not significant (Table 3). When we perform 5×5 median filtering twice ($2 \times$ **MF**) or five times ($5 \times$ **MF**) per warping step, the results are worse. Finally, removing the median filtering step (**w/o MF**) makes the computed flow significantly less accurate with larger outliers as shown in Table 3 and Fig. 4.

One interesting result with **HS** is that repeatedly applying median filtering (20 times) at every warping step improves the **HS** formulation and the improvement is statistically significant (**HS** $20 \times$ **MF** in Table 21).

Good Practices: Median filtering the intermediate flow results once after every warping iteration is the single most important implementation detail here; 5×5 is a good filter size.

4.6 Best Practices

Combining the analysis above into a single approach means modifying the baseline to use the slightly non-convex generalized Charbonnier and the spline-based bicubic interpolation. This leads to a statistically significant improvement over the baseline (Table 3, **Classic++**). This method is directly descended from **HS** and **BA**, yet updated with the current best optimization practices known to us. This simple method ranks 32th out of 73 methods in both EPE and AAE on the Middlebury test set at the writing of the paper (Sep. 2012). However, as we will see soon, this method is somehow not “simple”. Instead of the original objective, a different objective is being optimized with the median filtering step. The same is true for the reported results of both **HS** and **BA**.

	Sum	Venus	Dimetrodon	Hydrangea	RubberWhale	Grove2	Grove3	Urban2	Urban3
Bicubic-II	8.761	0.552	0.734	0.835	0.481	1.656	2.167	1.061	1.275
Deriv-warp	8.917	0.559	0.745	0.840	0.484	1.682	2.201	1.073	1.333
Warp-deriv	9.035	0.563	0.745	0.845	0.486	1.694	2.238	1.117	1.347

Table 4 Eq. (1) energy ($\times 10^6$) for the optical flow fields computed on the Middlebury *training set*, evaluated using spline-based bicubic interpolation [50]. Note the derivatives consistent with the interpolation method (**Bicubic-II**) produce the lowest energy solution.




Fig. 3 The median filtering is performed after every incremental warping step (i. e., once at every image pyramid level). The output of the median filtering is upsampled and used as the initial estimate for the next larger pyramid level.

	Sum	Venus	Dimetrodon	Hydrangea	RubberWhale	Grove2	Grove3	Urban2	Urban3
Classic-C	9.388	0.589	0.748	0.866	0.502	1.816	2.317	1.126	1.424
w/o GNC	9.689	0.593	0.750	0.870	0.506	1.845	2.518	1.142	1.465
w/o MF	8.044	0.517	0.701	0.668	0.449	1.418	1.830	1.066	1.395

Table 5 Eq. (1) energy ($\times 10^6$) for the optical flow fields computed on the Middlebury *training set*, evaluated using convolution-based bicubic interpolation [35]. Note that **Classic-C** uses graduated non-convexity (GNC), which reduces the energy, and median filtering, which increases it.




Fig. 4 Estimated flow fields on sequence “RubberWhale” using **Classic-C** with and without (**w/o MF**) the median filtering step. (a) (**w/ MF**) energy 502, 387, (b) (**w/o MF**) energy 449, 290, (c) color key [11]. The median filtering step helps reach a solution free from outliers but with a higher energy. The flow fields have been normalized by their maximum magnitude resulting in different contrasts. The outliers in the result without median filtering (b) make the flow appear lower contrast.

5 Models Underlying Median Filtering

Our analysis reveals the practical importance of median filtering during optimization. This effectively *denoises* the in-

termediate flow fields, preventing gross outliers, and making even non-robust methods like **HS** more robust. We ask whether there is a *principle* underlying this heuristic?

One interesting observation is that flow fields obtained with median filtering have substantially *higher* energy than those without (Table 5 and Fig. 4). If the median filter is helping to optimize the objective, it should lead to lower energies. Higher energies and more accurate estimates suggest that incorporating median filtering changes the objective function being optimized.

The insight that follows from this is that the median filtering heuristic is related to the minimization of an objective function that differs from the classical one. In particular the optimization of Eq. (1), with interleaved median filtering,




Fig. 5 From left to right, neighborhood structure for the center (red) pixel for the standard pairwise model, the unweighted non-local model, the unweighted non-local model with a larger neighborhood, and the weighted non-local model. The standard pairwise model connects a center pixel with its nearest neighbors, while the non-local term connects a pixel with many pixels in a large spatial neighborhood. By assigning larger weights (thicker red edges) to neighbors that are more likely to be on the same surface (blue circles), the weighted non-local model incorporates spatial scene structure information.

approximately minimizes

$$E(\mathbf{u}, \mathbf{v}) = \sum_{i,j} \left\{ \rho_D(I_1(i,j) - I_2(i + u_{i,j}, j + v_{i,j})) + \lambda[\rho_S(u_{i,j} - u_{i+1,j}) + \rho_S(u_{i,j} - u_{i,j+1}) + \rho_S(v_{i,j} - v_{i+1,j}) + \rho_S(v_{i,j} - v_{i,j+1})] \right\} + \lambda_N \sum_{i,j} \sum_{(i',j') \in \mathcal{N}_{i,j}} (|u_{i,j} - u_{i',j'}| + |v_{i,j} - v_{i',j'}|), \quad (2)$$

where $\mathcal{N}_{i,j}$ is the set of neighbors of pixel (i, j) in a possibly large area and λ_N is a scalar weight. The term in braces is the same as the flow energy from Eq. (1), while the last term is new. This *non-local* term [21, 28] imposes a particular smoothness assumption within a specified region of the flow field¹. Here we take this term to be a 5×5 rectangular region to match the size of the median filter in **Classic-C**. Figure 5 shows the neighborhood for the standard pairwise model and the non-local term.

It is usually difficult to directly optimize the objective (2) with a large spatial term. A common practice is to relax the objective with an auxiliary flow field as

$$E_A(\mathbf{u}, \mathbf{v}, \hat{\mathbf{u}}, \hat{\mathbf{v}}) = \sum_{i,j} \left\{ \rho_D(I_1(i,j) - I_2(i + u_{i,j}, j + v_{i,j})) + \lambda[\rho_S(u_{i,j} - u_{i+1,j}) + \rho_S(u_{i,j} - u_{i,j+1}) + \rho_S(v_{i,j} - v_{i+1,j}) + \rho_S(v_{i,j} - v_{i,j+1})] \right\} + \lambda_C(\|\mathbf{u} - \hat{\mathbf{u}}\|^2 + \|\mathbf{v} - \hat{\mathbf{v}}\|^2) + \lambda_N \sum_{i,j} \sum_{(i',j') \in \mathcal{N}_{i,j}} (|\hat{u}_{i,j} - \hat{u}_{i',j'}| + |\hat{v}_{i,j} - \hat{v}_{i',j'}|), \quad (3)$$

where $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ denote an auxiliary flow field and λ_C is a scalar weight. A third (coupling) term encourages $\hat{\mathbf{u}}, \hat{\mathbf{v}}$ and \mathbf{u}, \mathbf{v} to be the same (cf. [66, 75]). Here the notation implies a pixelwise sum of squared errors between the auxiliary and main flow fields.

¹ Bruhn *et al.* [20] also integrated information over a local region in a global method but did so for the data term.

The connection to median filtering (as a denoising method) derives from the fact that there is a direct relationship between the median and L1 minimization. Consider a simplified version of Eq. (3) with just the coupling and non-local terms, where

$$E(\hat{\mathbf{u}}) = \lambda_C \|\mathbf{u} - \hat{\mathbf{u}}\|^2 + \lambda_N \sum_{i,j} \sum_{(i',j') \in \mathcal{N}_{i,j}} |\hat{u}_{i,j} - \hat{u}_{i',j'}|. \quad (4)$$

While minimizing this is similar to median filtering \mathbf{u} , there are two differences. First, the non-local term minimizes the L1 distance between the central value and all flow values in its neighborhood except itself. Second, Eq. (4) incorporates information about the data term through the coupling equation; median filtering the flow ignores the data term.

The formal connection between Eq. (4) and median filtering² is provided by Li and Osher [40] who show that minimizing Eq. (4) is related to a different median computation

$$\hat{u}_{i,j}^{(k+1)} = \text{median}(\text{Neighbors}^{(k)} \cup \text{Data}) \quad (5)$$

where $\text{Neighbors}^{(k)} = \{\hat{u}_{i',j'}^{(k)}\}$ for $(i', j') \in \mathcal{N}_{i,j}$ and $\hat{\mathbf{u}}^{(0)} = \mathbf{u}$ as well as

$$\text{Data} = \{u_{i,j}, u_{i,j} \pm \frac{\lambda_N}{\lambda_C}, u_{i,j} \pm \frac{2\lambda_N}{\lambda_C} \dots, u_{i,j} \pm \frac{|\mathcal{N}_{i,j}|\lambda_N}{2\lambda_C}\},$$

where $|\mathcal{N}_{i,j}|$ denotes the (even) number of neighbors of (i, j) . Note that the set of “data” values is balanced with an equal number of elements on either side of the value $u_{i,j}$ and that information about the data term is included through $u_{i,j}$. Repeated application of Eq. (5) converges rapidly [40].

Observe that, as λ_N/λ_C increases, the weighted data values on either side of $u_{i,j}$ move away from the values of Neighbors and cancel each other out. As this happens, Eq. (5) approximates the median at the first iteration

$$\hat{u}_{i,j}^{(1)} \approx \text{median}(\text{Neighbors}^{(0)} \cup \{u_{i,j}\}). \quad (6)$$

Eq. (3) thus combines the original objective with an approximation to the median, the influence of which is controlled by λ_N/λ_C . Note in practice the weight λ_C on the coupling term is usually small or is steadily increased from small values [67, 75]. We optimize the new objective (3) by alternately minimizing

$$E_O(\mathbf{u}, \mathbf{v}) = \sum_{i,j} \left\{ \rho_D(I_1(i,j) - I_2(i + u_{i,j}, j + v_{i,j})) + \lambda[\rho_S(u_{i,j} - u_{i+1,j}) + \rho_S(u_{i,j} - u_{i,j+1}) + \rho_S(v_{i,j} - v_{i+1,j}) + \rho_S(v_{i,j} - v_{i,j+1})] \right\} + \lambda_C(\|\mathbf{u} - \hat{\mathbf{u}}\|^2 + \|\mathbf{v} - \hat{\mathbf{v}}\|^2) \quad (7)$$

² Hsiao *et al.* [32] established the connection in a slightly different way.

	Avg. EPE	significance	p-value
Classic-C	0.298	—	—
Classic-C-A	0.305	0	0.8125
Classic-C-A-noRep	0.309	0	0.5781
Classic-C-A-II	0.296	0	0.7188
Classic-C-A-CGD	0.305	0	0.5625

Table 6 Average end-point error (EPE) on the Middlebury *training set* is shown for the new model with alternating optimization (**Classic-C-A**). Please refer to Table 19 for the detailed EPE results on each training sequence.

and

$$E_M(\hat{\mathbf{u}}, \hat{\mathbf{v}}) = \lambda_C (\|\mathbf{u} - \hat{\mathbf{u}}\|^2 + \|\mathbf{v} - \hat{\mathbf{v}}\|^2) + \lambda_N \sum_{i,j} \sum_{(i',j') \in \mathcal{N}_{i,j}} (|\hat{u}_{i,j} - \hat{u}_{i',j'}| + |\hat{v}_{i,j} - \hat{v}_{i',j'}|). \quad (8)$$

We find that optimization of the coupled set of equations is superior in terms of EPE performance than optimization of the objective in Eq. (2).

The alternating optimization strategy first holds $\hat{\mathbf{u}}, \hat{\mathbf{v}}$ fixed and minimizes Eq. (7) w. r. t. \mathbf{u}, \mathbf{v} . Then, with \mathbf{u}, \mathbf{v} fixed, we minimize Eq. (8) w. r. t. $\hat{\mathbf{u}}, \hat{\mathbf{v}}$. Note that Eqs. (4) and (8) can be minimized by repeated application of Eq. (5); we use this approach with 5 iterations. We perform 10 steps of alternating optimizations at every pyramid level and change λ_C logarithmically from 10^{-4} to 10^2 . During the first and second GNC stages, we set \mathbf{u}, \mathbf{v} to be $\hat{\mathbf{u}}, \hat{\mathbf{v}}$ after every warping step (this replacement step helps reach solutions with lower energy and EPE than without performing this step; see **Classic-C-A-noRep** in Tables 6 and 7). In the end, we take $\hat{\mathbf{u}}, \hat{\mathbf{v}}$ as the final flow field estimate. The other parameters are $\lambda = 5, \lambda_N = 1$.

Alternately optimizing this new objective function (**Classic-C-A**) leads to similar results as the baseline **Classic-C** (Table 6). We also compare the energy of these solutions using the new objective and find the alternating optimization produces the lowest energy solutions, as shown in Table 7.

We find that approximately optimizing the new objective by changing λ_C logarithmically from 10^{-4} to 10^{-1} has slightly better EPE results but higher energy solutions (**Classic-C-A-II**). We also try replacing the absolute value by the Charbonnier penalty function and using the conjugate gradient descent method [6] to solve Eq. (4) but obtain results with slightly worse EPE performance and higher energy.

In summary, we show that the heuristic median filtering step in **Classic-C** can now be viewed as energy minimization of a new objective with a non-local term. The explicit formulation emphasizes the value of robustly integrating information over large neighborhoods and enables the improved model described below.

6 Improved Model

By formalizing the median filtering heuristic as an explicit objective function, we can find ways to improve it. While median filtering in a large neighborhood has advantages as we have seen, it also has problems. A neighborhood centered on a corner or thin structure is dominated by the surround and computing the median results in oversmoothing as illustrated in Fig. 1.

Examining the non-local term suggests a solution. For a given pixel, if we know which other pixels in the area belong to the same surface, we can weight them more highly. The modification to the objective function is achieved by introducing a weight into the non-local term [21, 28]:

$$\sum_{i,j} \sum_{(i',j') \in \mathcal{N}_{i,j}} w_{i,j}^{i',j'} (|\hat{u}_{i,j} - \hat{u}_{i',j'}| + |\hat{v}_{i,j} - \hat{v}_{i',j'}|), \quad (9)$$

where $w_{i,j}^{i',j'}$ represents how likely pixel i', j' is to belong to the same surface as i, j .

Of course, we do not know $w_{i,j}^{i',j'}$, but can approximate it. We draw ideas from [53, 71, 74] to define the weights according to their spatial distance, their color-value distance, and their occlusion state as

$$w_{i,j}^{i',j'} \propto \exp \left\{ \frac{-|i-i'|^2 + |j-j'|^2}{2\sigma_1^2} - \frac{|\mathbf{I}(i,j) - \mathbf{I}(i',j')|^2}{2\sigma_2^2 n_c} \right\} \frac{o(i',j')}{o(i,j)}, \quad (10)$$

where $\mathbf{I}(i,j)$ is the color vector in the Lab space, n_c is the number of color channels, $\sigma_1 = 7, \sigma_2 = 7$, and the occlusion variable $o(i,j)$ is calculated using Eq. (22) in [53] as

$$o(i,j) = \exp \left\{ -\frac{d^2(i,j)}{2\sigma_d^2} - \frac{(I(i,j) - I(i+u_{i,j}, j+v_{i,j}))^2}{2\sigma_e^2} \right\}, \quad (11)$$

where $d(i,j)$ is the one-sided divergence function, defined as

$$d(i,j) = \begin{cases} \text{div}(i,j), & \text{div}(i,j) < 0 \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

in which the flow divergence $\text{div}(i,j)$ is

$$\text{div}(i,j) = \frac{\partial}{\partial x} u(i,j) + \frac{\partial}{\partial y} v(i,j), \quad (13)$$

where $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$ are respectively the horizontal and vertical flow derivatives. The occlusion variable $o(i,j)$ is near zero for occluded pixels and near one for non-occluded pixels. We set the parameters in Eq. (11) as $\sigma_d = 0.3$ and $\sigma_e = 20$; this is the same as in [53]. Note that the occlusion state nonlinearly depends on the unknown flow field and we calculate the occlusion state using the latest flow estimate.

	Sum	Venus	Dimetrodon	Hydrangea	RubberWhale	Grove2	Grove3	Urban2	Urban3
Classic-C	13.013	0.817	0.903	1.202	0.674	2.166	3.144	1.954	2.153
Classic-C w/o MF	14.629	0.886	0.945	1.299	0.725	2.315	3.513	2.234	2.712
Classic-C-A	12.489	0.784	0.889	1.139	0.666	2.064	2.976	1.922	2.049
Classic-C-A-noRep	13.076	0.790	0.894	1.165	0.670	2.092	3.143	2.005	2.317
Classic-C-A-II	13.308	0.830	0.915	1.235	0.686	2.223	3.247	1.990	2.182
Classic-C-A-CGD	13.466	0.833	0.909	1.224	0.674	2.213	3.357	2.020	2.236

Table 7 Eq. (3) energy ($\times 10^6$) for the computed flow fields on the Middlebury *training set*. The alternating optimization strategy (**Classic-C-A**) produces the lower energy solutions than the median filtering heuristic.




Fig. 6 Neighbor weights of the proposed weighted non-local term at different positions in the “Army” sequence. We use color, spatial distance, and occlusion cues to determine whether the neighboring pixels are likely to belong to the same surface. Among these cues, color is the most powerful (see Table 8 and text for an evaluation of the cues).

Examples of such weights are shown for several 15×15 neighborhoods in Figure 6; bright values indicate higher weights. Note the neighborhood labeled **d**, corresponding to the rifle. Since pixels on the rifle are in the minority, an unweighted median oversmooths (**Classic++** in Fig. 1). The weighted term instead robustly estimates the motion using values on the rifle. A closely related piece of work is [51], which uses the intervening contour to define affinities among neighboring pixels for the local Lucas and Kanade [43] method. However it only uses this scheme to estimate motion for sparse points and then interpolates the dense flow field.

We approximately solve for \hat{u} (and similarly \hat{v}) using the following weighted median problem

$$\min_{\hat{u}_{i,j}} \sum_{(i',j') \in N_{i,j} \cup \{i,j\}} w_{i',j'} |\hat{u}_{i,j} - u_{i',j'}|, \quad (14)$$

using the formula (3.13) in [40] for all the pixels (**Classic+NL-Full**). Note if all the weights are equal, the solution is just the median. In practice, we can adopt a fast version (**Classic+NL**) without performance loss: Given a current estimate of the

	Avg. EPE	significance	p-value
Classic+NL	0.221	—	—
Classic+NL-Full	0.222	0	0.8203
Classic+NL-Fast	0.221	0	0.3125
RGB	0.240	1	0.0156
HSV	0.231	1	0.0312
LUV	0.226	0	0.5625
Gray	0.253	1	0.0078
w/o color	0.283	1	0.0156
w/o occ	0.226	0	0.1250
w/o spa	0.223	0	0.5625
$\sigma_2 = 5$	0.221	0	1.0000
$\sigma_2 = 10$	0.224	0	0.2500
$\lambda = 1$	0.236	0	0.1406
$\lambda = 9$	0.244	0	0.1016
11×11	0.223	0	0.5938
19×19	0.220	0	0.8750

Table 8 Average end-point error (EPE) on the Middlebury *training set* is shown for the the improved model and its variants. Please refer to Table 20 for the detailed results.

	Avg. Rank	Avg. EPE	Avg. EPE near boundary
Classic++	32.7	0.406	0.980
Classic++Gradient	33.5	0.430	1.042
Classic+NL	17.2	0.319	0.689
Classic+NL-Full	17.5	0.316	0.676

Table 9 Average end-point error (EPE) on the Middlebury *test set* for the **Classic++** model with two different preprocessing techniques and its improved model. Please refer to Table 15 for the detailed EPE results.

flow, we detect motion boundaries using a Sobel edge detector and dilate these edges with a 5×5 mask to obtain flow boundary regions. In these regions we use the weighting in Eq. (10) in a 15×15 neighborhood. In the non-boundary regions, we use equal weights in a 5×5 neighborhood to compute the median.

To further reduce the computation, we can adopt a two-stage GNC process and perform 3 warping steps per pyramid level. This fast version (**Classic+NL-Fast**) has nearly the same overall performance, with a slight decline in performance on the “Urban3” sequence, which has large motions; with an iterative warping scheme, large motions require more iterations.

Tables 8 and 9 show that the weighted non-local term (**Classic+NL**) improves the accuracy on both the training and the test sets, especially in the motion boundary regions. Note that the fine detail of the “rifle” is preserved in Figure 1(e). At the writing of this paper (Sep. 2012), **Classic+NL** ranks 13th in both AAE and EPE. Figures 7 and 8 show some of the results on the Middlebury dataset.




Fig. 7 Results on the Middlebury test set. Top to bottom: “Teddy”, “Wooden”, and “Grove”. **Classic+NL** uses information from the color image to detect and preserve fine motion details. Note that the ground truth visualization from the Middlebury website has been compressed and has lower quality than the actual ground truth.

We study some variants of the weighted non-local term (**Classic+NL**). Table 8 shows the importance of each term in determining the weight and influence of the parameter setting on the final results. Using different color spaces results in some performance decline. Using grayscale pixel values (**Gray**) or not using the static image information (**w/o color**) results in significant degradation in performance. Without occlusion (**w/o occ**) or spatial distance (**w/o spa**) cues does not degrade the performance significantly. The method is robust to the setting of σ_2 for the color cue and 5 and 10 perform similarly as the default 7. The default λ is 3, while 1 and 9 result in some loss in performance. We also study the maximum size of the neighborhood for the non-local term and find 11×11 gives similar performance while 19×19 is slightly better.

Closely-related work. Werlberger *et al.* [68] independently propose a non-local term for optical flow estimation and

the spatial term is similar to our non-local term. They use zero mean normalized cross correlation as the data term to deal with lighting changes. Their work is motivated by the success of the non-local regularization [21] in image restoration and stereo. Our work is inspired by the success of the heuristic median filtering step in flow estimation and we formalize the median filtering heuristic as a non-local regularization term. The use of the GPU and C++ makes their implementation faster than our implementation in MATLAB. **Classic+NL** has lower average EPE on the Middlebury test sequences; 0.319 versus 0.388 (cf. Table 15). Readers can visually compare the results of both methods on the Middlebury website.

Results on the MIT dataset. To test the robustness of these models on other data, we applied **HS**, **Classic-C**, and **Classic+NL** to sequences from the MIT dataset [42], and compared the estimated flow fields to the human labeled ground

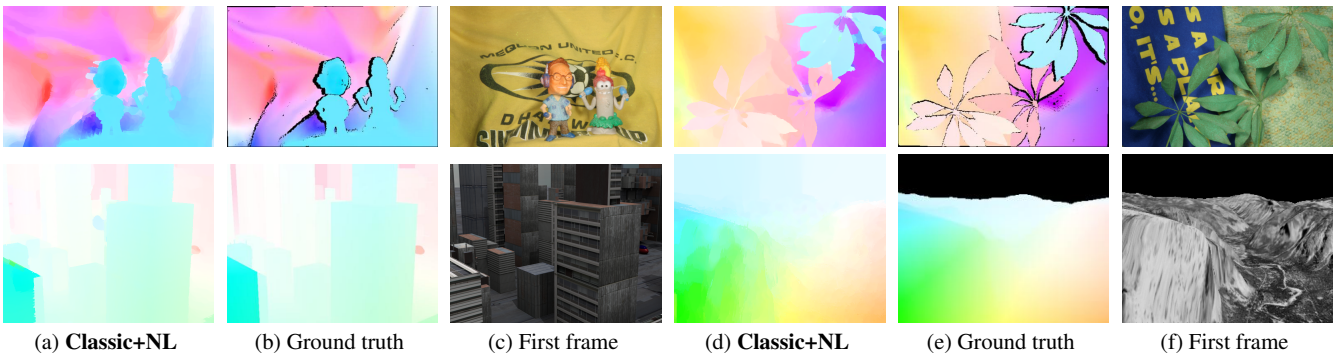


Fig. 8 Results on other Middlebury test sequences. Top left: “Mequon”; top right: “Schefflera”; bottom left: “Urban”; bottom right “Yosemite”.

truth. Note only five of the eight test sequences in [42] are available on-line; these are tested here.

Fig. 9 and Table 10 show the results on these sequences, which are very different in nature from the Middlebury set and include an outdoor scene as well as a scene of a fish tank. The results are compared with the CLG method [20] used in [42]. It is important to point out that the CLG method was tuned to obtain the optimal results on the test sequences. Our method had no such tuning and we used the same parameters as those used in all the other experiments. This suggests that training on the Middlebury data results in a method that generalizes to other sequences. The only place where this fails is on the “fish” sequence where there is transparent motion in a liquid medium; the statistics in this sequence are very different from the Middlebury training data.

	Average	Table	Hand	Toy	Fish	CameraMotion
CLG [20,42]	1.239	0.976	4.181	0.456	0.196	0.385
HS	2.129	1.740	6.108	0.620	1.309	0.869
Classic-C	1.345	1.064	3.428	0.482	1.061	0.690
Classic+NL	1.106	0.91	2.75	0.487	0.772	0.611

Table 10 Results on the MIT dataset [42]. Average end-point error (EPE). The CLG [20] method was tuned for each sequence [42].

Performance on MPI Sintel and KITTI datasets. We evaluate the methods above (corresponding to our publicly released code) on the MPI Sintel [23] and the KITTI [26] datasets using the default parameter settings in our conference paper [57]. As summarized in Tables 11 and 12, the conclusions contradict our findings reported above. On the MPI Sintel dataset, **HS** outperforms **Classic++**, which in turn outperforms **Classic+NL-fast**. The only consistent result is **Classic+NL**, which achieves the best performance. On the KITTI dataset, **HS** outperforms **Classic+NL**.

We ask how these datasets differ from both Middlebury and the MIT dataset? What could lead to these inconsistent conclusions? One answer surprisingly lies in the unequal width and height of the images.

Asymmetric pyramids for wide-aspect-ratio video. Our original implementation downsamples the image equally in the horizontal and vertical dimensions. The method automatically determines the number of pyramid levels using the smaller of the height and width of the input image. This scheme works well when the width-to-height ratio is close to 1, i. e., the Middlebury sequences. In contrast, the MPI Sintel images are 1024×436 and the KITTI images are around 1226×370 . The small vertical dimension limits the height of the pyramid, but we find that the large horizontal dimension means that the sequences contain very large horizontal motions. As a result, at the top level of the pyramid, the horizontal motions can be much larger than a pixel.

To address this we can use an unequal downsampling factor in each direction to ensure that the motion at the top pyramid level is small in both directions (or at least similar). For the MPI Sintel and KITTI data sets, we use a downsampling factor of 0.5 in the horizontal direction and determine the downsampling factor in the vertical direction and the pyramid level number, so that the size of top pyramid level is around 16×16 .

For MPI Sintel and KITTI this scheme results in a 7-level pyramid (instead of a 5-level pyramid in the standard symmetric scheme). This results in a significant improvement on both the the MPI Sintel and the KITTI data set, as summarized in Tables 11 and 12. We denote the method with the new asymmetric pyramid by adding an “A” at the end of the name.

On MPI Sintel, the results of the four methods are consistent with those on the Middlebury data set. Note that even **Classic++A** outperforms the previous **Classic+NL**. **Classic+NLA** outperforms **MDP-flow2** [73] on the final set, but not on the clean set. **MDP-flow2** uses feature matching to deal with fast moving objects. Feature matching tends to work well on the clean set, but not the final set due to motion and optical blur in the latter. Fig. 10 shows an example visual comparison between results using **Classic+NLA** and **Classic+NL**. The asymmetric pyramid leads to significant improvement in large areas that undergo large motions.

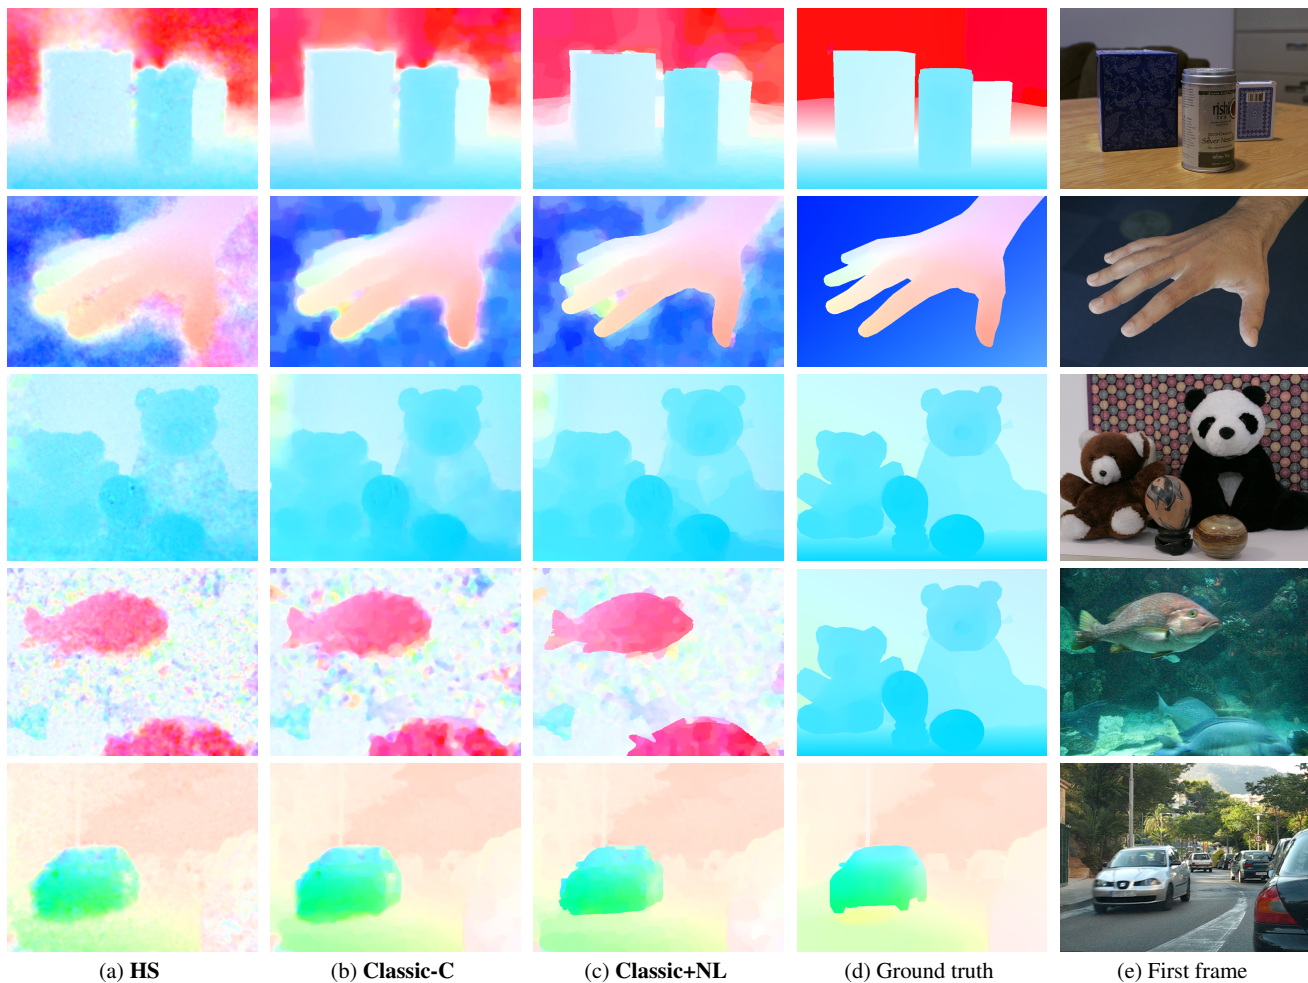


Fig. 9 Results on MIT sequences. Top to bottom: “Table”, “Hand”, “Toy”, “Fish”, and “CameraMotion”

Table 11 Average end-point error (EPE) on the MPI Sintel *training* set.

	Classic+NL	Classic+NL-fast	Classic++	HS
Final	7.998	8.806	8.601	8.386
Clean	6.035	7.445	6.765	7.238
	Classic+NLA	Classic+NL-FastA	Classic++A	HSA
Final	7.459	7.379	7.928	8.291
Clean	4.942	5.078	5.363	6.899

On the KITTI set, **Classic++A** performs best among all our tested methods, both in the training and the test sets. Note that the KITTI sequences have been collected on a moving vehicle in an urban environment. The flow fields tend to be smooth with few flow boundaries. The image-independent smoothness assumption in **Classic++A** is better suited to such data. Fig. 11 shows some results for **Classic+NL-FastA** and **Classic+NL-fast**; note the dramatic improvement resulting from the asymmetric pyramid.

It is important to note that, apart from the change of pyramid method, all other parameters remain the same and are trained using the Middlebury training sequences.

Table 12 Average end-point error (EPE) on the MPI Sintel *test* set. Only published results are shown, including **MDP-Flow2** [73], **FC-2Layers-FF** [61] and **AHOF**[69].

	Classic+NL	Classic+NL-fast	Classic++	HS
Final	9.153	10.088	9.959	9.610
Clean	7.961	9.129	8.721	8.739
	Classic+NLA	Classic+NL-FastA	Classic++A	HSA
Final	8.291	8.439	8.836	9.377
Clean	6.731	6.940	7.209	8.309
	MDP-Flow2	FC-2Layers-FF	AHOF	
Final	8.445	8.137	11.927	
Clean	5.837	6.781	12.642	

Computational time. Table 14 summarizes the running time of the evaluated methods on typical sequences from three different data sets in MATLAB on a 64-bit Linux desktop with 8GB memory. The additional cost from **HS** to **Classic++** comes from the GNC stage and the non-convex penalty function. The additional cost from **Classic++** to **Classic+NL** comes from the weighted median filtering step for detected motion boundaries. Applying the weighted median operation on all the pixels (**Classic+NL-Full**) increases




Fig. 10 Example results on MPI Sintel dataset. From top to bottom: first frame, second frame, results by **Classic+NL** (5-level), results by **Classic+NLA** (7-level), and ground truth. The asymmetric pyramid leads to a significant improvements in large regions undergoing large motion (head of the dragon on the left and background on the right). EPE results: “temple2” (left), 18.04 by **Classic+NL** (5-level) and 12.92 by **Classic+NLA** (7-level); “cave2” (right), 52.208 by **Classic+NL** (5-level) and 26.565 by **Classic+NLA** (7-level). Note that the estimated motion for fast-moving objects still contains large errors.

the running time by more than three times with little performance gain. Using fewer iterations (**Classic+NL-Fast**) can significantly reduce the computational cost with little performance loss, especially on sequences with small motion. Note that we solve the weighted median problem at each pixel individually and do not reuse the sorting results from neighboring pixels. Future work should consider reformulat-

ing the weighted median filtering so that a convolution-type operation can be used to reduce the computational cost.

Limitations. **Classic+NL** produces larger errors in occlusion regions on some sequences, such as “Schefflera” shown in Fig. 12. The classical flow formulation assumes that every pixel at the current frame has a corresponding pixel at the next frame. However, this assumption breaks down in




Fig. 11 Example results on the KITTI dataset. From top to bottom: first frame, second frame, results by **Classic+NL-fast** (5-level), results by **Classic+NL-FastA** (7-level), and ground truth for the non-occluded feature points. EPE results in non-occluded sparse feature points: “000002” (left), 12.124 by **Classic+NL-fast** (5-level) and 2.444 by **Classic+NL-FastA** (7-level); “000030” (right), 20.554 by **Classic+NL-fast** (5-level) and 0.615 by **Classic+NL-FastA** (7-level).

Table 13 Percentage of pixels with EPE larger than 3 pixels in non-occluded (**Out-Noc**) and all **Out-All** areas and average EPE in non-occluded (**Avg-Noc**) and all **Avg-All** areas on the KITTI *test* set.

Method	Out-Noc	Out-All	Avg-Noc	Avg-All
HS	19.92 %	28.86 %	5.8 px	11.7 px
Classic+NL	24.64 %	33.35 %	9.0 px	16.4 px
HSA	14.77 %	24.08 %	4.0 px	9.0 px
Classic+NL-FastA	12.42 %	22.27 %	3.2 px	7.8 px
Classic+NLA	10.60 %	20.66 %	2.8 px	7.2 px
Classic++A	10.16 %	20.29 %	2.6 px	7.1 px

Table 14 Running time (in minutes) for computing one optical flow field from an image pair from different benchmark dataset using different methods in MATLAB on a 64-bit Linux desktop with 8GB memory. Used sequences: 640×480 Urban from Middlebury, 1024×436 alley_1 from MPI Sintel, and 1226×370 training image 0 from KITTI.

	Middlebury	MPI Sintel	KITTI
HS	1.62	1.8	2.56
Classic++	5.83	7.2	8.48
Classic+NL	9.81	14	14.78
C+NL-fast	1.8	2.5	2.89
C+NL-full	26.7	29	42

regions of occlusion. Pixels that are occluded by some foreground objects in one frame do not have corresponding pixels in the next, resulting in large errors with classical formulations. In contrast, a layered model [64] may provide a principled way to reason about occlusions. The motion model developed in this paper has enabled a recent layered approach [59] to achieve a consistent improvement over the **Classic+NL** method, in particular near occlusion and motion boundary regions.

Small, fast moving objects also cause problems for the classical coarse-to-fine estimation used by **Classic+NL**, as shown in Fig. 10. The work by Brox and Malik [18] on large displacement optical flow has inspired recent work [24, 56, 73] to embed feature matching into the coarse-to-fine estimation framework. Chen *et al.* [24] show that, with proper initialization, **Classic+NL** can also handle large displacement optical flow on the Middlebury dataset.




Fig. 12 Occlusions are not explicitly modeled by **Classic+NL** and may cause problems in the estimated flow field. Dark pixels in the ground truth indicate occlusions.

7 Conclusions

When implemented using modern practices, classical optical flow formulations can produce fairly competitive results on existing datasets. To understand the techniques that help such basic formulations work well, we quantitatively studied various aspects of flow approaches from the literature, including their implementation details. Among the best practices, we found that using median filtering to denoise the flow after every warping step is key to improving accuracy, but that this increases the energy of the final result. Exploiting connections between median filtering and L1-based denoising, we showed that algorithms relying on a median filtering step are approximately optimizing a different objective that regularizes the flow field over a large spatial neighborhood. Understanding this enables us to design and optimize improved models that weight the neighbors adaptively in an extended image region. The MATLAB code is publicly available [1].

There has been much debate about whether methods that perform well on Middlebury will generalize to other sequences. Here we tuned the parameters of the method on the Middlebury training set and tested on Middlebury, MIT HAMA, MPI Sintel, and KITTI. The conclusions on the Middlebury dataset are consistent with those on the MIT HAMA dataset. The one significant difference we found between Middlebury and the MPI Sintel and KITTI datasets was the aspect ratio of the images. This allowed us to make a change to the method by introducing a novel asymmetric image pyramid that downsamples more rapidly in the horizontal direction than the vertical direction. With only this change we found that our conclusions on Middlebury hold for MPI Sintel as well. The KITTI data set is somewhat different in nature and seems to favor methods with more spatial smoothing. As a result, the image-independent **Classic++**, which produces more smooth flow fields, performs slightly better than the image-dependent **Classic+NL**, with its sharp boundaries. It is open whether these conclusions will hold for data tak-

en under totally different conditions, such as medical images. While the results on Middlebury generalize surprisingly well, we suspect that training the parameters for a specific dataset will improve results further.

Acknowledgments. DS and MJB were supported by a gift from Intel Corp. and NSF CRCNS award IIS-0904875. We thank the CVPR and IJCV reviewers for constructive comments, especially the connection between our original “area” term and non-local regularization and test on additional datasets, P. Yadollahpour for his early work on implementing the HS method, S. Zuffi for suggesting the color version of the non-local term, J. Wulff for running some experiments, T. Brox, A. Wedel, and M. Werlberger for clarifying details about their papers, and D. Scharstein for maintaining the online optical flow benchmark.

References

1. <http://www.cs.brown.edu/people/dqsun>.
2. <http://gpu4vision.icg.tugraz.at>.
3. <http://lmb.informatik.uni-freiburg.de/resources/software.php>.
4. <http://people.csail.mit.edu/celiu/OpticalFlow/>.
5. <http://www.cse.cuhk.edu.hk/~leojia/projects/flow/>.
6. <http://www.gaussianprocess.org/gpml/code/matlab/util/minimize.m>.
7. Y. Adato, T. Zickler, and O. Ben-Shahar. A polar representation of motion and implications for optical flow. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1145–1152, 2011.
8. E. H. Adelson, C. H. Anderson, J. R. Bergen, P. J. Burt, and J. M. Ogden. Pyramid methods in image processing. *RCA Engineer*, 29(6):33–41, November 1984.
9. L. Álvarez, C. A. Castaño-Moraga, M. García, K. Krissian, L. Mazorra, A. Salgado, and J. Sánchez. Symmetric optical flow. In *EUROCAST*, pages 676–683, 2007.
10. A. Bab-Hadiashar and D. Suter. Robust optic flow computation. *International Journal of Computer Vision*, 29(1):59–77, August 1998.
11. S. Baker, D. Scharstein, J. Lewis, S. Roth, M. J. Black, and R. Szeliski. A database and evaluation methodology for optical flow. In *IEEE International Conference on Computer Vision*, pages 1–8, 2007.
12. J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, February 1994.
13. J. Bergen, P. Anandan, K. Hanna, and R. Hingorani. Hierarchical model-based motion estimation. In *European Conference on Computer Vision*, volume 588, pages 237–252, 1992.
14. M. Black and A. Jepson. Estimating optical-flow in segmented images using variable-order parametric models with local deformations. *IEEE Transaction on Pattern Analysis Machine Intelligence*, 18(10):972–986, October 1996.
15. M. J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer Vision and Image Understanding*, 63:75–104, 1996.
16. A. Blake and A. Zisserman. *Visual Reconstruction*. The MIT Press, Cambridge, Massachusetts, 1987.
17. T. Brox, A. Bruhn, N. Papenberger, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *European Conference on Computer Vision*, pages 25–36, 2004.

Table 15 Models. Average end-point error (EPE) on the Middlebury optical flow benchmark (*test set*). The ranking information was determined at the writing of the paper (Sep. 2012).

	Rank	Average	Army	Mequon	Schefflera	Wooden	Grove	Urban	Yosemite	Teddy
HS	49.0	0.501	0.12	0.25	0.45	0.24	0.95	0.83	0.24	0.93
Classic-C	34.8	0.408	0.10	0.23	0.45	0.20	0.88	0.47	0.16	0.77
Classic-L	42.7	0.530	0.10	0.24	0.47	0.21	0.92	1.23	0.20	0.87
HS-brightness	N/A	0.759	0.21	0.89	1.13	0.42	0.93	0.70	0.18	1.61
Classic-C-brightness	N/A	0.726	0.39	0.95	1.12	0.42	0.87	0.48	0.13	1.45
Classic-L-brightness	N/A	0.603	0.17	0.64	0.84	0.32	0.90	0.48	0.13	1.34
HS [58]	66.2	0.872	0.22	0.61	1.01	0.78	1.26	1.43	0.16	1.51
BA (Classic-L) [58]	59.6	0.746	0.18	0.58	0.95	0.49	1.08	1.43	0.15	1.11
Adaptive [66]	28.5	0.401	0.09	0.23	0.54	0.18	0.88	0.50	0.14	0.65
Complementary OF [77]	31.6	0.485	0.10	0.20	0.35	0.19	0.87	1.46	0.11	0.60
NL-TV-NCC [68]	23.5	0.388	0.10	0.22	0.35	0.15	0.79	0.78	0.16	0.55
Classic++	32.7	0.406	0.09	0.23	0.43	0.20	0.87	0.47	0.17	0.79
Classic++Gradient	33.5	0.430	0.08	0.17	0.49	0.21	0.94	0.55	0.17	0.83
Classic+NL	17.2	0.319	0.08	0.22	0.29	0.15	0.64	0.52	0.16	0.49
Classic+NL-Full	17.5	0.316	0.08	0.24	0.28	0.15	0.63	0.49	0.16	0.50

Table 16 Models and pre-processing. Average end-point error (EPE) on the Middlebury *training set* for the classical model and different penalty functions. By default, the input sequences were preprocessed using ROF texture decomposition; “brightness” means no preprocessing is performed. The statistical significance is tested using the Wilcoxon signed rank test between each method and the baseline (**Classic-C**).

	Average	Venus	Dimetrodon	Hydrangea	RubberWhale	Grove2	Grove3	Urban2	Urban3	signif.	p-value
Classic-C	0.298	0.281	0.152	0.165	0.093	0.158	0.627	0.348	0.562	—	—
Classic-C-brightness	0.288	0.268	0.166	0.215	0.134	0.146	0.584	0.352	0.437	0	0.9453
HS	0.384	0.337	0.219	0.189	0.118	0.204	0.688	0.463	0.853	1	0.0078
HS-brightness	0.387	0.335	0.226	0.252	0.154	0.185	0.639	0.564	0.743	1	0.0078
Classic-L	0.319	0.294	0.193	0.175	0.095	0.166	0.648	0.374	0.604	1	0.0078
Classic-L-brightness	0.325	0.292	0.207	0.274	0.145	0.158	0.588	0.451	0.484	0	0.2969

Table 17 Pre-Processing. Average end-point error (EPE) on the Middlebury *training set* for the baseline method (**Classic-C**) using different pre-processing techniques. The regularization weight λ parameter was tuned for each method to achieve optimal performance. The statistical significance is tested using the Wilcoxon signed rank test between each method and the baseline (**Classic-C**).

	Average	Venus	Dimetrodon	Hydrangea	RubberWhale	Grove2	Grove3	Urban2	Urban3	signif.	p-value
Classic-C	0.298	0.281	0.152	0.165	0.093	0.158	0.627	0.348	0.562	—	—
Gradient	0.305	0.288	0.141	0.167	0.092	0.165	0.614	0.385	0.588	0	0.4609
Gaussian	0.281	0.268	0.146	0.226	0.141	0.137	0.582	0.335	0.413	0	0.5469
Gaussian + Dx + Dy	0.290	0.280	0.126	0.174	0.105	0.154	0.588	0.470	0.420	0	0.6406
Dx + Dy	0.301	0.286	0.122	0.166	0.099	0.161	0.616	0.443	0.518	0	1.0000
Sobel edge[63]	0.417	0.334	0.149	0.184	0.130	0.194	0.757	0.451	1.135	1	0.0156
Laplacian [38]	0.430	0.374	0.170	0.176	0.096	0.175	0.756	0.464	1.232	1	0.0078
Laplacian 1:1	0.301	0.296	0.179	0.193	0.109	0.157	0.606	0.349	0.520	0	0.6641
Texture 4:1	0.286	0.271	0.159	0.175	0.100	0.154	0.587	0.349	0.490	0	0.5312
Unnormalized texture	0.298	0.279	0.152	0.166	0.092	0.158	0.623	0.348	0.563	0	0.3750

Table 18 Model and Methods. Average end-point error (EPE) on the Middlebury *training set* for the baseline model (**Classic-C**) using different algorithm and modeling choices. The statistical significance is tested using the Wilcoxon signed rank test between each method and the baseline (**Classic-C**).

	Average	Venus	Dimetrodon	Hydrangea	RubberWhale	Grove2	Grove3	Urban2	Urban3	signif.	p-value
Classic-C	0.298	0.281	0.152	0.165	0.093	0.158	0.627	0.348	0.562	—	—
3 warping steps	0.304	0.283	0.122	0.163	0.095	0.150	0.622	0.357	0.644	0	0.9688
Down-0.5	0.298	0.280	0.152	0.166	0.092	0.158	0.626	0.349	0.562	0	1.0000
Down-0.95	0.298	0.281	0.151	0.168	0.099	0.165	0.661	0.339	0.523	0	0.9375
w/o GNC	0.354	0.303	0.160	0.171	0.105	0.183	0.835	0.316	0.759	0	0.1094
Bilinear	0.302	0.284	0.144	0.167	0.099	0.160	0.637	0.363	0.563	0	0.1016
w/o TAVG	0.306	0.288	0.149	0.167	0.093	0.163	0.647	0.345	0.593	0	0.1562
Central	0.300	0.272	0.156	0.169	0.092	0.159	0.608	0.349	0.597	0	0.7266
7-point [20]	0.302	0.282	0.168	0.171	0.091	0.163	0.601	0.360	0.584	0	0.3125
Deriv-warp	0.297	0.283	0.153	0.165	0.092	0.159	0.636	0.333	0.552	0	0.9531
Bicubic-II	0.290	0.276	0.132	0.152	0.083	0.142	0.624	0.338	0.571	1	0.0391
Deriv-warp-II	0.287	0.264	0.155	0.152	0.085	0.145	0.616	0.333	0.546	1	0.0156
Warp-deriv-II	0.288	0.267	0.155	0.151	0.085	0.147	0.630	0.328	0.542	1	0.0391
C-L ($\lambda = 0.6$)	0.303	0.290	0.158	0.171	0.094	0.158	0.611	0.367	0.579	0	0.1562
L-C ($\lambda = 2$)	0.306	0.281	0.174	0.173	0.096	0.164	0.662	0.343	0.557	0	0.1562
GC-0.45 ($\lambda = 3$)	0.292	0.280	0.145	0.165	0.092	0.154	0.612	0.340	0.546	1	0.0156
GC-0.25 ($\lambda = 0.7$)	0.298	0.283	0.128	0.169	0.094	0.150	0.617	0.353	0.594	0	1.0000
MF 3×3	0.305	0.287	0.155	0.168	0.094	0.162	0.616	0.372	0.583	0	0.1016
MF 7×7	0.305	0.281	0.152	0.173	0.095	0.174	0.676	0.330	0.557	0	0.5625
$2 \times$ MF	0.300	0.279	0.152	0.167	0.093	0.163	0.650	0.339	0.555	0	1.0000
$5 \times$ MF	0.305	0.278	0.152	0.171	0.093	0.172	0.682	0.329	0.561	0	0.6875
w/o MF	0.352	0.307	0.168	0.199	0.113	0.217	0.705	0.423	0.684	1	0.0078
Classic++	0.285	0.271	0.128	0.153	0.081	0.139	0.614	0.336	0.555	1	0.0078

Table 19 Average end-point error (EPE) on the Middlebury *training set* for the proposed new objective with the non-local term and alternating optimization (**Classic-C-A**) and its improved models. The statistical significance is tested using the Wilcoxon signed rank test between each method and the baseline (**Classic-C**).

	Average	Venus	Dimetrodon	Hydrangea	RubberWhale	Grove2	Grove3	Urban2	Urban3	signif.	p-value
Classic-C	0.298	0.281	0.152	0.165	0.093	0.158	0.627	0.348	0.562	—	—
Classic-C-A	0.305	0.281	0.140	0.159	0.092	0.167	0.676	0.334	0.594	0	0.8125
Classic-C-A-noRep	0.309	0.279	0.139	0.161	0.093	0.157	0.653	0.370	0.619	0	0.5781
Classic-C-A-II	0.296	0.278	0.153	0.166	0.091	0.168	0.656	0.329	0.531	0	0.7188
Classic-C-A-CGD	0.305	0.281	0.148	0.161	0.093	0.159	0.697	0.344	0.560	0	0.5625

Table 20 Average end-point error (EPE) on the Middlebury *training set* for the proposed new objective with the weighted non-local term and its variants. The statistical significance is tested using the Wilcoxon signed rank test between each method and the baseline (**Classic+NL**).

	Average	Venus	Dimetrodon	Hydrangea	RubberWhale	Grove2	Grove3	Urban2	Urban3	signif.	p-value
Classic+NL	0.221	0.238	0.131	0.152	0.073	0.103	0.468	0.220	0.384	—	—
Classic+NL-Full	0.222	0.252	0.135	0.156	0.074	0.097	0.469	0.214	0.382	0	0.8203
Classic+NL-Fast	0.221	0.233	0.117	0.151	0.076	0.098	0.464	0.210	0.421	0	0.3125
RGB	0.240	0.243	0.131	0.155	0.081	0.109	0.501	0.236	0.468	1	0.0156
HSV	0.231	0.245	0.131	0.152	0.074	0.110	0.492	0.222	0.424	1	0.0312
LUV	0.226	0.241	0.131	0.149	0.074	0.104	0.460	0.223	0.427	0	0.5625
Gray	0.253	0.253	0.133	0.158	0.086	0.125	0.547	0.242	0.479	1	0.0078
w/o color	0.283	0.258	0.128	0.157	0.087	0.155	0.633	0.303	0.543	1	0.0156
w/o occ	0.226	0.243	0.131	0.152	0.073	0.103	0.488	0.230	0.386	0	0.1250
w/o spa	0.223	0.237	0.132	0.154	0.073	0.102	0.475	0.213	0.398	0	0.5625
$\sigma_2 = 5$	0.221	0.240	0.131	0.151	0.073	0.104	0.466	0.208	0.392	0	1.0000
$\sigma_2 = 10$	0.224	0.238	0.132	0.153	0.073	0.102	0.485	0.228	0.384	0	0.2500
$\lambda = 1$	0.236	0.245	0.151	0.164	0.080	0.120	0.430	0.243	0.459	0	0.1406
$\lambda = 9$	0.244	0.249	0.137	0.160	0.091	0.111	0.577	0.201	0.426	0	0.1016
11×11	0.223	0.240	0.131	0.151	0.074	0.103	0.451	0.234	0.397	0	0.5938
19×19	0.220	0.238	0.132	0.154	0.073	0.103	0.470	0.210	0.384	0	0.8750

Table 21 Additional results for HS. Average end-point error (EPE) on the Middlebury *training set*. The statistical significance is tested using the Wilcoxon signed rank test between each method and **HS**.

	Average	Venus	Dimetrodon	Hydrangea	RubberWhale	Grove2	Grove3	Urban2	Urban3	signif.	p-value
HS	0.384	0.337	0.219	0.189	0.118	0.204	0.688	0.463	0.853	—	—
HS 20 × MF	0.365	0.299	0.214	0.184	0.104	0.196	0.699	0.431	0.792	1	0.0469

18. T. Brox and J. Malik. Large displacement optical flow: Descriptor matching in variational motion estimation. *IEEE Transaction on Pattern Analysis Machine Intelligence*, 33(3):500–513, March 2011.
19. A. Bruhn and J. Weickert. Towards ultimate motion estimation: Combining highest accuracy with real-time performance. In *ICCV*, pages 749–755, 2005.
20. A. Bruhn, J. Weickert, and C. Schnörr. Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. *International Journal of Computer Vision*, 61(3):211–231, February 2005.
21. A. Buades, B. Coll, and J. Morel. A non-local algorithm for image denoising. In *IEEE International Conference on Computer Vision and Pattern Recognition*, volume 2, pages 60–65, 2005.
22. P. J. Burt, C. Yen, and X. Xu. Local correlation measures for motion analysis: A comparative study. *Processings of IEEE Pattern Recognition and Image Processing*, pages 269–274, 1982.
23. D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conference on Computer Vision*, IV, pages 611–625, 2012.
24. Z. Chen, H. Jin, Z. Lin, S. Cohen, and Y. Wu. Large displacement optical flow from nearest neighbor fields. In *IEEE International Conference on Computer Vision and Pattern Recognition*, 2013. to appear.
25. Z. Chen, Y. Wu, and J. Wang. Decomposing and regularizing sparse/non-sparse components for motion field estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1176–1183, 2012.
26. A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, Providence, USA, 2012.
27. S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transaction on Pattern Analysis Machine Intelligence*, 6(6):721–741, November 1984.
28. G. Gilboa and S. Osher. Nonlocal operators with applications to image processing. *SIAM Multiscale Modeling and Simulation*, 7:1005–1028, 2008.
29. F. Glaer, G. Reynolds, and P. Anandan. Scene matching by hierarchical correlation. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 432–441, 1983.
30. B. Horn. *Robot Vision*. MIT Press, 1986.
31. B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 16(1-3):185–203, August 1981.
32. I. Hsiao, A. Rangarajan, and G. Gindi. A new convex edge-preserving median prior with applications to tomography. *IEEE Transactions on Medical Imaging*, 22(5):580–585, May 2003.
33. A. Humayun, O. Mac Aodha, and G. J. Brostow. Learning to Find Occlusion Regions. In *IEEE International Conference on Computer Vision and Pattern Recognition*, number 2161–216, 2011.
34. K. Jia, X. Wang, and X. Tang. Optical flow estimation using learned sparse model. In *IEEE International Conference on Computer Vision*, pages 2391–2398, 2011.
35. R. G. Keys. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 29(6):1153–1160, December 1981.
36. P. Krähenbühl and V. Koltun. Efficient nonlocal regularization for optical flow. In *European Conference on Computer Vision*, volume I, pages 356–369, 2012.
37. C. Lei and Y.-H. Yang. Optical flow estimation on coarse-to-fine region-trees using discrete optimization. In *IEEE International Conference on Computer Vision*, pages 1562–1569, 2009.
38. V. Lempitsky, S. Roth, and C. Rother. FusionFlow: Discrete-continuous optimization for optical flow estimation. In *IEEE In-*

- ternational Conference on Computer Vision and Pattern Recognition*, 2008.
39. V. Lempitsky, C. Roth, S. Roth, and A. Blake. Fusion moves for Markov random field optimization. *IEEE Transaction on Pattern Analysis Machine Intelligence*, 32(8):1392–1405, August 2010.
 40. Y. Li and S. Osher. A new median formula with applications to PDE based denoising. *Communications in Mathematical Sciences*, 7(3):741–753, September 2009.
 41. D. Lin and J. Fisher. Low level vision via switchable Markov random fields. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 2432–2439, 2012.
 42. C. Liu, W. T. Freeman, E. H. Adelson, and Y. Weiss. Human-assisted motion annotation. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
 43. B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *International Joint Conferences on Artificial Intelligence*, pages 674–679, 1981.
 44. O. Mac Aodha, G. J. Brostow, and M. Pollefeys. Segmenting video into classes of algorithm-suitability. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1778–1785, 2010.
 45. D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W.H. Freeman, New York, NY, USA, 1982.
 46. H.-H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transaction on Pattern Analysis Machine Intelligence*, 8(5):565–593, September 1986.
 47. M. Nikolova. Model distortions in Bayesian MAP reconstruction. *AIMS J. on Inverse Problems and Imaging*, 1:399–422, 2007.
 48. T. Nir, A. M. Bruckstein, and R. Kimmel. Over-parameterized variational optical flow. *International Journal of Computer Vision*, 76(2):205–216, 2008.
 49. Y. Niu, W. Feng, and F. Liu. Enabling warping on stereoscopic images. *ACM Trans. Graph.*, 31(6):183, 2012.
 50. W. H. Press, W. T. Vetterling, S. A. Teukolsky, and B. P. Flannery. *Numerical Recipes in C++: The art of scientific computing*. Cambridge University Press, New York, NY, USA, 2002.
 51. X. Ren. Local grouping for optical flow. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
 52. L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, November 1992.
 53. P. Sand and S. Teller. Particle video: Long-range motion estimation using point trajectories. *International Journal of Computer Vision*, 80(1):72–91, October 2008.
 54. U. Schmidt, Q. Gao, and S. Roth. A generative perspective on MRFs in low-level vision. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1751–1758, 2010.
 55. D. Shulman and J.-Y. Herve. Regularization of discontinuous flow fields. In *Workshop on Visual Motion*, pages 81–86, 1989.
 56. F. Steinbrücker, T. Pock, and D. Cremers. Large displacement optical flow computation without warping. In *IEEE International Conference on Computer Vision*, pages 1609–1614, 2009.
 57. D. Sun, S. Roth, and M. J. Black. Secrets of optical flow estimation and their principles. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 2432–2439, 2010.
 58. D. Sun, S. Roth, J. P. Lewis, and M. J. Black. Learning optical flow. In *European Conference on Computer Vision*, pages 83–97, 2008.
 59. D. Sun, E. B. Sudderth, and M. J. Black. Layered image motion with explicit occlusions, temporal consistency, and depth ordering. In *Advances in Neural Information Processing Systems*, pages 2226–2234, 2010.
 60. D. Sun, E. B. Sudderth, and M. J. Black. Layered segmentation and optical flow estimation over time. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1768–1775, 2012.
 61. D. Sun, J. Wulff, E. B. Sudderth, H. Pfister, and M. J. Black. A fully-connected layered model of foreground and background flow. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1768–1775, 2013.
 62. R. Szeliski. *Computer Vision: Algorithms and Applications*. Springer-Verlag New York, Inc., 2010.
 63. T. Vaudrey and R. Klette. Residual images remove illumination artifacts! In *Pattern Recognition (Proceedings of DAGM)*, pages 472–481, Berlin, Heidelberg, 2009. Springer-Verlag.
 64. J. Y. A. Wang and E. H. Adelson. Representing moving images with layers. *IEEE Transactions on Image Processing*, 3(5):625–638, September 1994.
 65. A. Wedel, T. Pock, J. Braun, U. Franke, and D. Cremers. Duality TV-L1 flow with fundamental matrix prior. In *Image and Vision Computing New Zealand*, 2008.
 66. A. Wedel, T. Pock, and D. Cremers. Structure- and motion-adaptive regularization for high accuracy optic flow. In *IEEE International Conference on Computer Vision*, pages 1663–1668, 2009.
 67. A. Wedel, T. Pock, C. Zach, D. Cremers, and H. Bischof. An improved algorithm for TV-L1 optical flow. In *Dagstuhl Motion Workshop*, pages 23–45, 2008.
 68. M. Werlberger, T. Pock, and H. Bischof. Motion estimation with non-local total variation regularization. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 2464–2471, 2010.
 69. M. Werlberger, W. Trobin, T. Pock, A. Wedel, D. Cremers, and H. Bischof. Anisotropic Huber-L1 optical flow. In *Proceedings of the British Machine Vision Conference*, pages 108.1–108.11, 2009.
 70. F. Wilcoxon. Individual Comparisons by Ranking Methods. *Biometrics Bulletin*, 1(6):80–83, December 1945.
 71. J. Xiao, H. Cheng, H. Sawhney, C. Rao, and M. Isnardi. Bilateral filtering-based optical flow estimation with occlusion detection. In *European Conference on Computer Vision*, volume I, pages 211–224, 2006.
 72. L. Xu, J. Chen, and J. Jia. A segmentation based variational model for accurate optical flow estimation. In *European Conference on Computer Vision*, volume I, pages 671–684, 2008.
 73. L. Xu, J. Jia, and Y. Matsushita. Motion detail preserving optical flow estimation. *IEEE Transaction on Pattern Analysis Machine Intelligence*, 34(9):1744–1757, September 2012.
 74. K. Yoon and I. Kweon. Adaptive support-weight approach for correspondence search. *IEEE Transaction on Pattern Analysis Machine Intelligence*, 28(4):650–656, April 2006.
 75. C. Zach, T. Pock, and H. Bischof. A duality based approach for realtime TV-L1 optical flow. In *Pattern Recognition (Proceedings of DAGM)*, pages 214–223, 2007.
 76. H. Zimmer, A. Bruhn, and J. Weickert. Optic flow in harmony. *International Journal of Computer Vision*, 93(3):368–388, 2011.
 77. H. Zimmer, A. Bruhn, J. Weickert, L. Valgaerts, A. Salgado, B. Rosenhahn, and H.-P. Seidel. Complementary optic flow. In *Energy Minimization Methods in Computer Vision and Pattern*, pages 207–220, 2009.
 78. C. Zitnick, N. Jojic, and S. B. Kang. Consistent segmentation for optical flow estimation. In *IEEE International Conference on Computer Vision*, volume 2, pages 1308–1315, 2005.