

## Overview

### Goals:

- What makes modern optical flow techniques accurate and why?
- How can we use such insights to improve flow techniques further?

**Secrets:** Quantitative analysis of current practices in optical flow estimation starting from a simple, classical formulation

**Principles:** Formalization of the heuristic median filtering step as an unweighted non-local term added to the original objective

**Improved model:** Introduce weighted non-local term that uses color, flow, and occlusion information to better preserve motion details

**MATLAB code:** <http://www.cs.brown.edu/~dqsun/>

## Secrets

What makes optical flow accurate?

Classical formulation:

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

Standard data constancy term

Smoothness term: Pairwise Markov random field (MRF) with 4 neighbors

Modern implementation:

- Preprocessing: Rudin-Osher-Fatemi (ROF) structure texture decomposition [3]
- Standard incremental multi-resolution technique, 10 warping steps per level
- Graduate non-convexity (GNC) for non-quadratic penalty
- Bicubic interpolation to warp image and its derivatives
- 5-point image derivative filter
- Temporal averaging of image derivatives
- 5X5 median filtering of intermediate flow field per warping step [3]

How good is it?

	Middlebury test set	Avg. Rank	Avg. EPE	
Current implementation	<b>Classic-Charbonnier</b>	14.9	0.408	$\rho(x) = \sqrt{x^2 + \epsilon^2}$
	<b>HS (Horn &amp; Schunk)</b>	24.6	0.501	$\rho(x) = x^2$
	<b>Classic-Lorentizan (BA)</b>	19.8	0.530	$\rho(x) = \log(1 + \frac{x^2}{2\sigma^2})$
Previous implementation	<b>HS</b>	35.1	0.872	
	<b>BA (Black &amp; Anandan)</b>	30.9	0.746	
Previous state of the art	<b>Adaptive</b>	11.5	0.401	
	<b>Complementary OF</b>	10.1	0.485	
For later comparison	<b>Classic-C-brightness</b>	N/A	0.726	

Table 1. **Models.** Average end-point error (EPE) on Middlebury test set by Classical formulation with different penalty functions.

What is important?

Approach: Change one property of **Classic-C** at a time, compare avg. EPE, and test statistical significance

## Secrets uncovered

• **Pre-processing:** Some kind of image filtering is useful but simple gradient constancy is as good as more sophisticated texture decomposition; overfitting is more severe for brightness constancy

• **Interpolation and image derivatives:** Bicubic interpolation is slightly better than bilinear but not significantly; *spline-based bicubic interpolation* is consistently better than convolution-based (MATLAB built-in); removing temporal averaging of image derivatives, central difference filter, and 7-point derivative filter reduce accuracy, but not significantly

• **Coarse-to-fine estimation and GNC:** Pyramid downsampling factor does not matter for the convex penalty and 0.5 is fine; graduated non-convexity (GNC) helps even the convex robust Charbonnier penalty

• **Penalty function:** Less robust Charbonnier is better than Lorentzian; a slightly more robust penalty (*generalized Charbonnier*  $\rho(x) = (x^2 + \epsilon^2)^a$ ,  $a = 0.45$ ) is better still

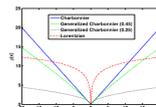


Figure 1. Different penalty functions for the spatial terms.

• **Median filtering:** Median filtering the intermediate flow field is the single most important secret; 5x5 is a good filter size

Wilcoxon signed rank test between each variant and baseline Classic-C

	Middlebury training set	Avg. EPE	signif.	p-value
Baseline	Classic-C	0.298	-	-
Preprocessing	Brightness constancy	0.288	0	0.9453
	Gradient constancy	0.305	0	0.4609
	Bilinear interpolation	0.302	0	0.1016
Interpolation and image derivatives	Central difference filter	0.300	0	0.7266
	7-point derivative filter	0.302	0	0.3125
	<b>Spline-based bicubic interpolation</b>	<b>0.290</b>	<b>1</b>	<b>0.0391</b>
Coarse-to-fine estimation and GNC	No temporal average of derivatives	0.306	0	0.1562
	Downsampling factor 0.5	0.298	0	1.0000
	3 warping steps per level	0.304	0	0.9688
Penalty function	No graduated non-convexity (GNC)	0.354	0	0.1094
	<b>Generalized Charbonnier-0.45</b>	<b>0.292</b>	<b>1</b>	<b>0.0156</b>
	Generalized Charbonnier-0.25	0.298	0	1.0000
Median filtering	Median filter size 3 X 3	0.305	0	0.1016
	Median filter size 7 X 7	0.305	0	0.5625
	Median filtering twice	0.300	0	1.0000
Best practices	<b>No median filtering</b>	<b>0.352</b>	<b>1</b>	<b>0.0078</b>
	<b>Classic++</b>	<b>0.285</b>	<b>1</b>	<b>0.0078</b>

Table 2. **Models and Methods.** Average end-point error (EPE) on Middlebury training set for **Classic-C** and its variants.

**Best practices (Classic++):** Modify baseline Classic-C to use the slightly non-convex generalized Charbonnier and spline-based bicubic interpolation. This method is directly descended from HS and BA, yet updated with current best practices known to us. It ranks 8<sup>th</sup> in EPE on the public Middlebury test set.

## Principles

Median filtering leads to lower EPE, but higher energy solutions!



With MF: Energy **502,387**, EPE 0.093 Without MF: Energy **449,290**, EPE 0.113

Figure 2. Estimated flow fields on "RubberWhale" by **Classic-C**.

What is being minimized?

**Observation:** Median filtering can be posed as L1 energy minimization [1]. Replace median filter with minimization of this objective function:

$$E(\hat{\mathbf{u}}) = \lambda_2 \|\mathbf{u} - \hat{\mathbf{u}}\|^2 + \sum_{i,j} \sum_{(i',j') \in N_{i,j}} |\hat{u}_{i,j} - \hat{u}_{i',j'}|$$

**New objective function:** Non-local term robustly integrates information over a **large spatial neighborhood**

$$E_A(\mathbf{u}, \mathbf{v}, \hat{\mathbf{v}}) = \sum_{i,j} \{ \rho_D(I_1(i, j) - I_2(i + u_{i,j}, j + v_{i,j})) + \lambda [\rho_S(u_{i,j} - u_{i+1,j}) + \rho_S(u_{i,j} - u_{i,j+1}) + \rho_S(v_{i,j} - v_{i+1,j}) + \rho_S(v_{i,j} - v_{i,j+1})] \}$$

$$+ \lambda_2 (\|\mathbf{u} - \hat{\mathbf{u}}\|^2 + \|\mathbf{v} - \hat{\mathbf{v}}\|^2) \quad \text{Coupling term}$$

$$+ \sum_{i,j} \sum_{(i',j') \in N_{i,j}} \lambda_3 (|\hat{u}_{i,j} - \hat{u}_{i',j'}| + |\hat{v}_{i,j} - \hat{v}_{i',j'}|) \quad \text{Non-local term}$$

**Optimization:** Alternate optimization between coupled classical and non-local terms

Alternating optimization (**Classic-C-A**) leads to similar performance

Middlebury training set	Avg. EPE	Significance	p-value
<b>Classic-C</b>	0.298	-	-
<b>Classic-C-A</b>	0.305	0	0.8125

Table 3. Average EPE on training set for the new objective with alternating optimization.



Table 4. Screen shot of the Middlebury optical flow benchmark (June 2010).

Average endpoint error	rank	Army (Hidden texture)				Mequon (Hidden texture)				Schefflera (Hidden texture)				Wooden (Hidden texture)				Grove (Synthetic)				Urban (Synthetic)				Yosemite (Synthetic)				Teddy (Stereo)			
		all	disc	untxt	all	disc	untxt	all	disc	untxt	all	disc	untxt	all	disc	untxt	all	disc	untxt	all	disc	untxt	all	disc	untxt	all	disc	untxt					
<b>Classic+NL [38]</b>	5.8	0.08	0.23	0.07	0.22	0.74	0.18	0.29	0.65	0.19	0.15	0.73	0.09	0.64	0.93	0.47	0.52	1.12	0.33	0.16	0.21	0.13	0.29	0.49	0.98	0.74							
MDP-Flow [30]	6.4	0.09	0.25	0.08	0.19	0.54	0.18	0.24	0.55	0.20	0.16	0.91	0.09	0.74	1.06	0.61	0.46	1.02	0.35	0.12	0.14	0.11	0.17	0.78	1.68	0.97							
NL-TV-NCC [28]	7.4	0.10	0.26	0.08	0.22	0.72	0.15	0.35	0.85	0.16	0.15	0.70	0.09	0.79	1.16	0.51	0.78	1.38	0.48	0.16	0.21	0.15	0.26	0.55	1.16	0.55							
Layer+dense [37]	8.1	0.09	0.28	0.08	0.22	0.74	0.19	0.25	0.58	0.21	0.17	0.92	0.09	0.87	1.17	0.94	0.35	1.05	0.31	0.16	0.21	0.13	0.28	0.55	1.17	0.79							
Complementary OF [24]	9.1	0.10	0.26	0.09	0.20	0.70	0.14	0.35	0.85	0.16	0.19	1.05	0.10	0.87	1.25	0.71	1.46	1.61	0.73	0.11	0.12	0.21	0.18	0.60	1.37	0.80							
Adaptive [23]	10.2	0.09	0.26	0.06	0.23	0.78	0.18	0.54	1.19	0.21	0.18	0.91	0.10	0.88	1.25	0.73	0.50	1.28	0.31	0.14	0.14	0.16	0.22	0.65	1.37	0.79							
Adapt-Window [36]	11.5	0.10	0.24	0.09	0.19	0.59	0.15	0.27	0.64	0.17	0.18	0.82	0.11	0.74	1.07	0.56	1.78	1.73	0.95	0.22	0.38	0.16	0.45	0.70	1.11	0.88							
<b>Classic++ [39]</b>	12.2	0.09	0.25	0.07	0.23	0.78	0.19	0.43	1.21	0.22	0.20	1.11	0.17	0.87	1.30	0.66	0.47	1.62	0.33	0.17	0.26	0.14	0.32	0.79	1.64	0.92							

## Improved model

**Problem:** Unweighted non-local term (median filtering) destroys fine structures that differ from the majority of neighbors

**Solution:** Weight neighbors adaptively according to color, flow, and occlusion state information

$$\sum_{i,j} \sum_{(i',j') \in N_{i,j}} w_{i,j,i',j'} (|\hat{u}_{i,j} - \hat{u}_{i',j'}| + |\hat{v}_{i,j} - \hat{v}_{i',j'}|)$$

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

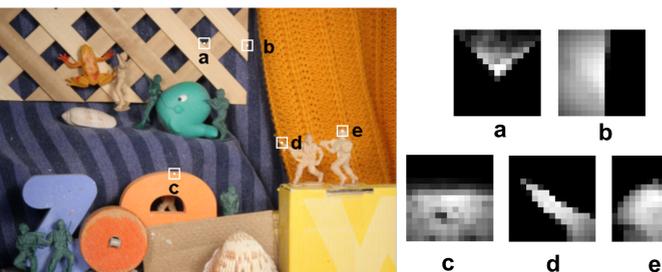


Figure 3. Neighbor weights of the proposed non-local term at different positions in the "Army" sequence.



Figure 4. Median filtering over-smooths the rifle in the "Army" sequence while the proposed non-local term preserves the details.

## Conclusions

- Classical formulations competitive with current practices
- Median filtering is key to accuracy, but increases energy
- Formalize median filtering as a non-local term that integrates information over a large spatial neighborhood
- Weighting neighbors adaptively preserves motion details
- **MATLAB code:** <http://www.cs.brown.edu/~dqsun/>

## References

- [1] Li and Osher. A new median formula with applications to PDE based denoising. Commun. Math. Sci., 7(3):741-753, 2009.
- [2] Sand and Teller. Particle video: Long-range motion estimation using point trajectories. IJCV, 2008.
- [3] Wedal et al. An improved algorithm for TV-L1 optical flow. Dagstuhl Motion Workshop, 2008.
- [4] Yoon and Kweon. Adaptive support-weight approach for correspondence search. PAMI, 2006.