



Interactive Media Systems Group
Software Technology & Interactive Systems
Vienna University of Technology



PatchMatch Stereo – Stereo Matching with Slanted Support Windows

Michael Bleyer (TU Vienna)

Christoph Rhemann (TU Vienna)


Carsten Rother (Microsoft Research Cambridge)

BMVC 2011

Local Stereo Algorithms



Left image



**Compare color
values within
search windows**

Right image

Local Stereo Algorithms

Implicit smoothness assumption:

- All pixels within support window have the **same integer-valued** disparity

Problems:

- Sub-pixel disparities
- Slanted surfaces
- Disparity discontinuities

Local Stereo Algorithms

Implicit smoothness assumption:

- All pixels within support window have the **same integer-valued** disparity

Problems:

- Sub-pixel disparities
- Slanted surfaces
- **Disparity discontinuities**

We use adaptive
Support Weights
[Yoon, CVPR05]

Local Stereo Algorithms

Implicit smoothness assumption:

- All pixels within support window have the **same integer-valued** disparity


Problems:

- **Sub-pixel disparities**
- **Slanted surfaces**
- **Disparity discontinuities**


Focus of our work

We use adaptive
Support Weights
[Yoon, CVPR05]

Fronto-Par. Windows Matched at Integer Disparities



Left image



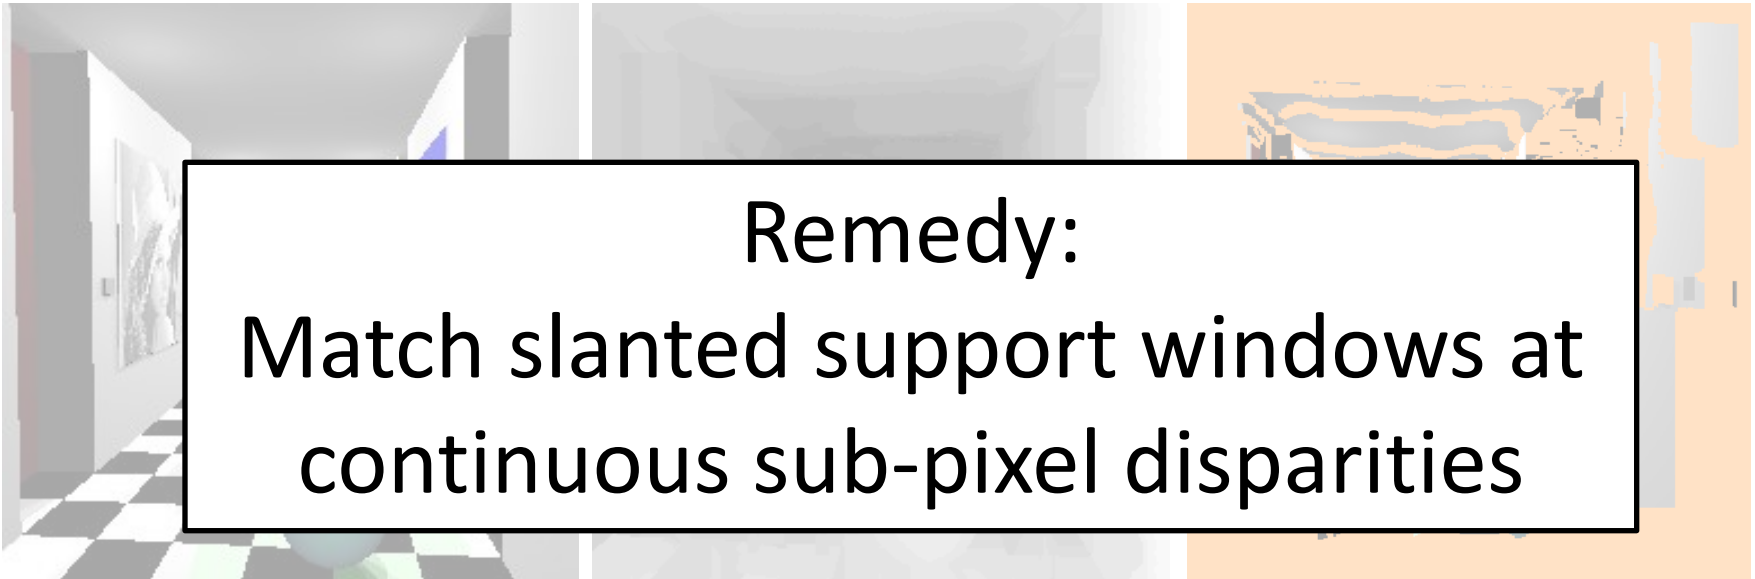
Disparity map



3D reconstruction

Fronto-Par. Windows Matched at Integer Disparities

Remedy:
Match slanted support windows at
continuous sub-pixel disparities




Left image

Disparity map


3D reconstruction

Slanted Support Windows




camera

Slanted Support Windows




Our Slanted Windows



Left image



Disparity map



3D reconstruction

Our Slanted Windows



The challenge:

“How can we find the correct slanted plane at each pixel?”

(infinite number of candidate planes)

Previous Work

- Sub-pixel precision and slanted windows via extending label space:
 - Include fronto-parallel planes at sub-pixel disparities
 - Include slanted planes (plane sweeping)
- Problems
 - Each additional label adds extra computational complexity
 - Algorithm fails if correct plane not part of label space

Previous Work

- S
e
• Our algorithm:
 - Continuous optimization over the set of all planes
 - No quantization needed
 - Based on PatchMatch [Barnes et al., SIGGRAPH09]
- P
• Algorithm fails if correct plane not part of label space


Our Algorithm

Basic Idea

- Relatively large regions of pixels can be modeled by approximately the same plane.
- Initialize each pixel with a random plane
- At least one pixel of a region should carry a plane close to correct one:
 - Many guesses (one per pixel)
- A single good guess is enough - it will be propagated to neighboring pixels



Left image – Sawtooth
(Middlebury)



Ground truth disparities





Image consists of 3 planes -
~80.000 guesses for yellow plane

Algorithm

- After random initialization process pixels in the following order:
 - Even iterations:



- Odd iterations:
 - Reverse order
- We run 3 iterations

Algorithm

- After random initialization process pixels in the following order:

Run the following pipeline for each pixel p :

1. Spatial propagation
2. View propagation
3. Temporal propagation
4. Plane refinement


- Odd iterations:
 - Reverse order
- We run 3 iterations

Spatial Propagation

- Look at the current pixel p 's spatial neighbors.
- Check if assigning p to a spatial neighbor's plane leads to lower aggregated costs.



Left image –
Reindeer
(Middlebury)




Left and right disparity maps (intermediate step of iteration 1)

View Propagation

- Find all pixels that have p as their matching point according to their current disparity.
- Check if plane of a matching point improves costs.



Left image –
Reindeer
(Middlebury)




Left and right disparity maps (intermediate step of iteration 1)

Temporal Propagation

- Applicable if operating on video sequences
- Find planes at p 's coordinates in the previous and consecutive frames.
- Check if planes improve current costs.


Plane Refinement

- Take p 's current plane f_p .




Plane Refinement

- Take p 's current plane f_p .
- Distance δ defines the maximum allowed change of f_p .



Plane Refinement


- Take p 's current plane f_p .
- Distance $delta$ defines the maximum allowed change of f_p .
- Loop
 - Compute a random plane f_p^* within $delta$



Plane Refinement

- Take p 's current plane f_p .
- Distance $delta$ defines the maximum allowed change of f_p .
- Loop
 - Compute a random plane f_p^* within $delta$
 - If f_p^* improves costs
 $f_p := f_p^*$


No improvement;
Dismiss f_p^*



Plane Refinement


- Take p 's current plane f_p .
- Distance $delta$ defines the maximum allowed change of f_p .
- Loop
 - Compute a random plane f_p^* within $delta$
 - If f_p^* improves costs
 $f_p := f_p^*$

No improvement;
Dismiss f_p^*




Plane Refinement

- Take p 's current plane f_p .
- Distance $delta$ defines the maximum allowed change of f_p .
- Loop
 - Compute a random plane f_p^* within $delta$
 - If f_p^* improves costs
$$f_p := f_p^*$$
 - $delta := delta / 2$ (exponential reduction)
 - Break if $delta < eps$




Plane Refinement

- Take p 's current plane f_p .
- Distance $delta$ defines the maximum allowed change of f_p .
- Loop
 - Compute a random plane f_p^* within $delta$
 - If f_p^* improves costs
 $f_p := f_p^*$
 - $delta := delta / 2$ (exponential reduction)
 - Break if $delta < eps$



Plane Refinement

- Take p 's current plane f_p .
- Distance $delta$ defines the maximum allowed change of f_p .
- Loop
 - Compute a random plane f_p^* within $delta$
 - If f_p^* improves costs
 - $f_p := f_p^*$
 - $delta := delta / 2$ (exponential reduction)
 - Break if $delta < eps$



Plane Refinement

- Take p 's current plane f_p .
- Distance $delta$ defines the maximum allowed change of f_p .
- Loop
 - Compute a random plane f_p^* within $delta$
 - If f_p^* improves costs
 $f_p := f_p^*$
 - $delta := delta / 2$ (exponential reduction)
 - Break if $delta < eps$

Improvement;
Accept f_p^*

disparities



We show the different steps of our
algorithm.

Global Methods

- Our slanted windows can be used as a data term for global methods (see paper).

Results

Evaluation

- We implement two competitors in our PatchMatch framework.
 - Competitor 1:
 - **Fronto-parallel** windows matched at **integer** disparities
 - Competitor 2:
 - **Fronto-parallel** windows matched at **sub-pixel** disparities
- Testbed:
 - Middlebury

Error threshold 1 (Middlebury default)

Algorithm	thresh.	Rank	Tsukuba			Venus			Teddy			Cones		
			nocc	all	disc	nocc	all	disc	nocc	all	disc	nocc	all	disc
Slanted Support	1.0	11	2.09 ₅₇	2.33 ₄₃	9.31 ₅₄	0.21 ₂₀	0.39 ₁₅	2.62 ₂₆	2.99₁	8.16 ₈	9.62 ₂	2.47 ₄	7.80 ₈	7.11 ₆
Fronto-Par. Sub-Pixel	1.0	22	2.13 ₅₇	2.32 ₄₃	9.92 ₆₀	0.25 ₂₇	0.42 ₁₆	2.76 ₂₇	5.52 ₁₇	10.5 ₁₇	14.2 ₁₅	3.23 ₂₇	8.33 ₁₇	8.20 ₁₈
Fronto-Par. Integer	1.0	32	2.23 ₅₈	2.44 ₄₄	9.18 ₅₃	0.25 ₂₅	0.41 ₁₆	2.21 ₁₆	6.73 ₃₄	12.2 ₃₄	16.9 ₃₈	3.71 ₃₅	8.90 ₂₇	9.31 ₃₅

Ranking in the Middlebury online table

Error threshold 1 (Middlebury default)

- Fronto-parallel windows at integer disparities (Comp. 1)
- Rank **32** of **110** Middlebury submissions

Algorithm	thresh.	Rank												
Slanted Support	1.0	11												
Fronto-Par. Sub-Pixel	1.0	2	2.13 ₅₇	2.32 ₄₃	9.92 ₆₀	0.25 ₂₇	0.42 ₁₆	2.76 ₂₇	5.52 ₁₇	10.5 ₁₇	14.2 ₁₅	3.23 ₂₇	8.33 ₁₇	8.20 ₁₈
Fronto-Par. Integer	1.0	32	2.23 ₅₈	2.44 ₄₄	9.18 ₅₃	0.25 ₂₅	0.41 ₁₆	2.21 ₁₆	6.73 ₃₄	12.2 ₃₄	16.9 ₃₈	3.71 ₃₅	8.90 ₂₇	9.31 ₃₅

Ranking in the Middlebury online table



Left image – Teddy set
 (Middlebury)



Disparity map



Error map (black = wrong)

Error threshold 1 (Middlebury default)


- Fronto-parallel windows at sub-pixel disp. (Comp. 2)
- Rank **22** of **110** Middlebury submissions

Algorithm	thresh.	Rank												
Slanted Support	1.0	1	2.09 ₅₇	2.33 ₄₃	9.31 ₅₄	0.21 ₂₀	0.39 ₁₅	2.62 ₂₆	2.99 ₁	8.16 ₈	9.62 ₅	2.47 ₄	7.80 ₈	7.11 ₆
Fronto-Par. Sub-Pixel	1.0	22	2.13 ₅₇	2.32 ₄₃	9.92 ₆₀	0.25 ₂₇	0.42 ₁₆	2.76 ₂₇	5.52 ₁₇	10.5 ₁₇	14.2 ₁₅	3.23 ₂₇	8.33 ₁₇	8.20 ₁₈
Fronto-Par. Integer	1.0	32	2.23 ₅₈	2.44 ₄₄	9.18 ₅₃	0.25 ₂₅	0.41 ₁₆	2.21 ₁₆	6.73 ₃₄	12.2 ₃₄	16.9 ₃₈	3.71 ₃₅	8.90 ₂₇	9.31 ₃₅

Ranking in the Middlebury online table



Left image – Teddy set
(Middlebury)



Disparity map



Error map (black = wrong)

Error threshold

- Our slanted windows
- Rank **11** of **110** Middlebury submissions
- Best-performing local method

Algorithm	thresh.	Rank	nocc	all	disc	nocc	all	disc	nocc	all	disc	nocc	all	disc
Slanted Support	1.0	11	2.09 ₅₇	2.33 ₄₃	9.31 ₅₄	0.21 ₂₀	0.39 ₁₅	2.62 ₂₆	2.99₁	8.16 ₈	9.62₂	2.47 ₄	7.80 ₈	7.11 ₆
Fronto-Par. Sub-Pixel	1.0	22	2.13 ₅₇	2.32 ₄₃	9.92 ₆₀	0.25 ₂₇	0.42 ₁₆	2.76 ₂₇	5.52 ₁₇	10.5 ₁₇	14.2 ₁₅	3.23 ₂₇	8.33 ₁₇	8.20 ₁₈
Fronto-Par. Integer	1.0	32	2.23 ₅₈	2.44 ₄₄	9.18 ₅₃	0.25 ₂₅	0.41 ₁₆	2.21 ₁₆	6.73 ₃₄	12.2 ₃₄	16.9 ₃₈	3.71 ₃₅	8.90 ₂₇	9.31 ₃₅

Ranking in the Middlebury online table



Left image – Teddy set
 (Middlebury)



Disparity map




Error map (black = wrong)

Error threshold

- Our slanted windows
- Rank **11** of **110** Middlebury submissions
- Best-performing local method

Algorithm	thresh.	Rank	nocc	all	disc	nocc	all	disc	nocc	all	disc	nocc	all	disc
Slanted Support	1.0	11	2.09 ₅₇	2.33 ₄₃	9.31 ₅₄	0.21 ₂₀	0.39 ₁₅	2.62 ₂₆	2.99₁	8.16 ₈	9.62 ₂	2.47 ₄	7.80 ₈	7.11 ₆
Fronto-Par. Sub-Pixel	1.0	22	2.13 ₅₇	2.32 ₄₃	9.92 ₆₀	0.25 ₂₇	0.42 ₁₆	2.76 ₂₇	5.52 ₁₇	10.5 ₁₇	14.2 ₁₅	3.23 ₂₇	8.33 ₁₇	8.20 ₁₈
Fronto-Par. Integr.									34	12.2 ₃₄	16.9 ₃₈	3.71 ₃₅	8.90 ₂₇	9.31 ₃₅

- Rank #1 on Teddy in non-occluded regions



Left image – Teddy set
 (Middlebury)

Disparity map

Error map (black = wrong)


Error threshold 0.5 (Sub-Pixel Evaluation)

Algorithm	thresh.	Rank	Tsukuba			Venus			Teddy			Cones		
			nocc	all	disc	nocc	all	disc	nocc	all	disc	nocc	all	disc
Slanted Support	0.5	2	15.0 ₃₇	15.4 ₃₆	20.3 ₄₅	1.00 ₄	1.34 ₄	7.75 ₇	5.66₁	11.8₂	16.5₁	3.80 ₁	10.2 ₁	10.2 ₁
Fronto-Par. Sub-Pixel	0.5	4	14.1 ₃₆	14.4 ₃₂	19.2 ₃₅	1.73 ₇	2.15 ₇	7.84 ₇	9.63 ₆	16.2 ₈	22.6 ₈	7.08 ₂₄	12.6 ₁₆	13.2 ₁₁
Fronto-Par. Integer	0.5	41												

- Our slanted windows
- Rank **2** of **110** Middlebury submissions



Left image – Teddy set
(Middlebury)



Disparity map



Error map (black = wrong)


Error threshold 0.5 (Sub-Pixel Evaluation)

Algorithm	thresh.	Rank	Tsukuba			Venus			Teddy			Cones		
			nocc	all	disc	nocc	all	disc	nocc	all	disc	nocc	all	disc
Slanted Support	0.5	2	15.0 ₃₇	15.4 ₃₆	20.3 ₄₅	1.00 ₄	1.34 ₄	7.75 ₇	5.66₁	11.8 ₂	16.5₁	3.80₁	10.2₁	10.2₁
Fronto-Par. Sub-Pixel	0.5	4	14.1 ₃₆	14.4 ₃₂	19.2 ₃₅	1.73 ₇	2.15 ₇	7.84 ₇	9.63 ₆	16.2 ₈	22.6 ₈	7.08 ₂₄	12.6 ₁₆	13.2 ₁₁
Fronto-Par. Pixel									29.1 ₃₇			11.6 ₅₃	16.9 ₄₂	18.3 ₄₁

- First rank on almost all error measurements on the complex Teddy and Cones sets




Left image – Teddy set
 (Middlebury)



Disparity map



Error map (black = wrong)



Computational Speed

- Approximately 1 minute per Middlebury pair (CPU)
- Runtime independent of the number of labels (disparities)
 - Well suited for optical flow computation (future work)
- Low memory requirements:
 - We only need to hold the current matching costs and a plane at each pixel in memory
 - => We can do high-resolution stereo

Image of Graz (2048x2048 pixels)



Image of Graz (2048x2048 pixels)




Image of Graz (2048x2048 pixels)



Crop of left image

Crop of right image



Left disparity map
(before left/right check)

Playroom Set (image courtesy Daniel Scharstein)




Playroom Set (image courtesy Daniel Scharstein)



Playroom Set (image courtesy Daniel Scharstein)



Playroom Set (image courtesy Daniel Scharstein)



Conclusions

- Local stereo algorithm uses slanted support windows
- An ideal algorithm to find the slanted windows is PatchMatch
- High quality results on complex images and videos at reasonable runtimes

Future Work

- Optical flow
- Improve computational speed (GPU implementation – real time)