Semi-Global Stereo Matching with Surface Orientation Priors

Daniel Scharstein — Middlebury College
Tatsunori Taniai — RIKEN AIP
Sudipta N. Sinha — Microsoft Research

3DV 2017

Contributions

- SGM's fronto parallel surface assumption makes it fail on untextured slanted surfaces
- our extension (SGM-P) utilizes precomputed surface orientation priors to reduce such errors

Advantages

- minimal runtime overhead
- easy to add to existing SGM implementation

Semi-Global Stereo Matching (SGM): Review

- Approximates 2D MRF using 1D optimization for 8 cardinal directions

$$E(D) = \sum_{\mathbf{p}} C_{\mathbf{p}}(d_{\mathbf{p}}) + \sum_{\mathbf{p},\mathbf{q}\in\mathcal{N}} V(d_{\mathbf{p}}, d_{\mathbf{q}})$$

- First order smoothness; fronto-parallel bias

$$V(d, d') = \begin{cases} 0 & \text{if } d = d' \\ P_1 & \text{if } |d - d'| = 1 \\ P_2 & \text{if } |d - d'| \geq 2 \end{cases}$$

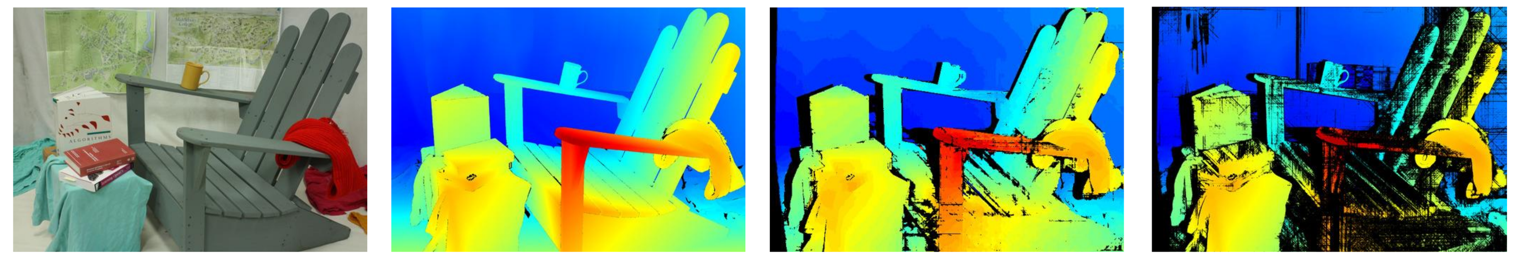

Left Image | Disparity (GT) | Disparity (Qrtr.) | Disparity (Full)

Where do we get priors?

- Matched features + triangulation
- Matched features + plane fitting
- Low-res matching + plane fitting
- Semantic prediction
- Manhattan-world assumptions
- GT oracles

Input | SGM-EPi | SGM-GS

Main Idea: If we knew the surface slants, we could replace fronto parallel bias with bias parallel to surface. But, how would we do it in a discrete optimization setting?

The solution is to *rasterize a disparity surface* at an arbitrary depth and adjust *V(d, d')* to follow the discrete steps.

2D Orientation Prior:

$$V_S(d_{\mathbf{p}}, d'_{\mathbf{p}}) = V(d_{\mathbf{p}} + j_{\mathbf{p}}, d'_{\mathbf{p}})$$

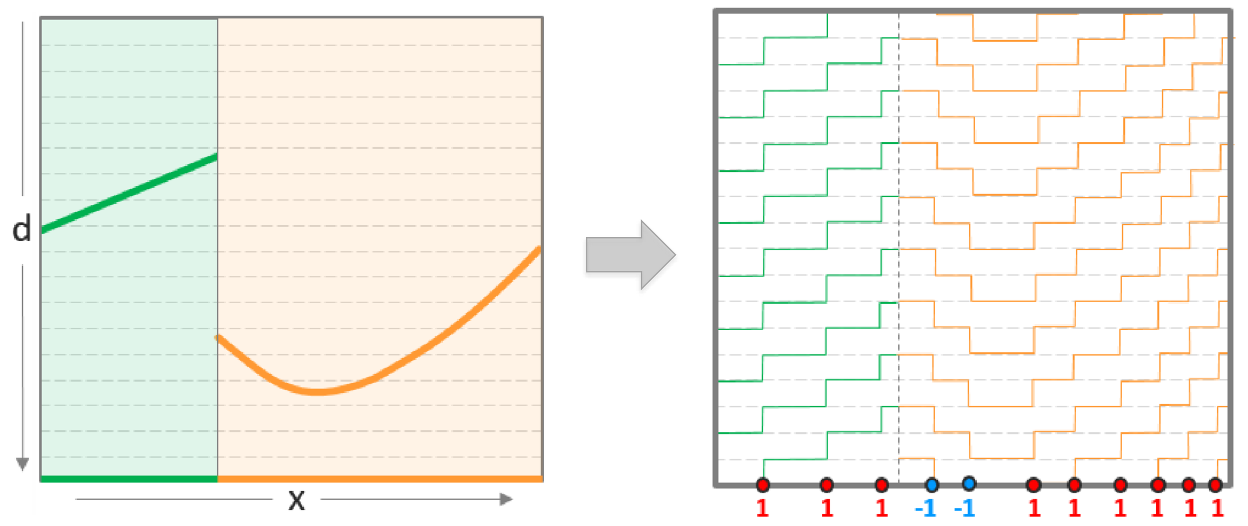

offset depends on pixel location

$$j_{\mathbf{p}} = \{-1, 0, +1\}$$

3D Orientation Prior:

$$V_S(d_{\mathbf{p}}, d'_{\mathbf{p}}) = V(d_{\mathbf{p}} + j_{\mathbf{p}}(d_{\mathbf{p}}), d'_{\mathbf{p}})$$

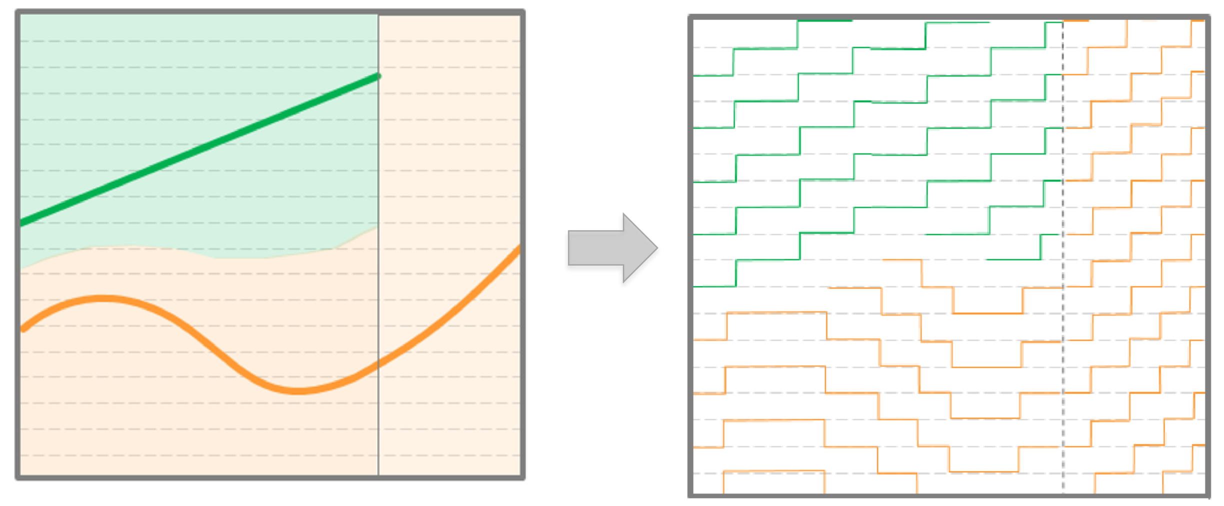

offset depends on pixel location and disparity

Surface Normal Priors

- We derive the exact relationship between surface normals in scene coordinates and disparity surface orientations
- Requires 3D orientation prior; disparity steps no longer aligned vertically

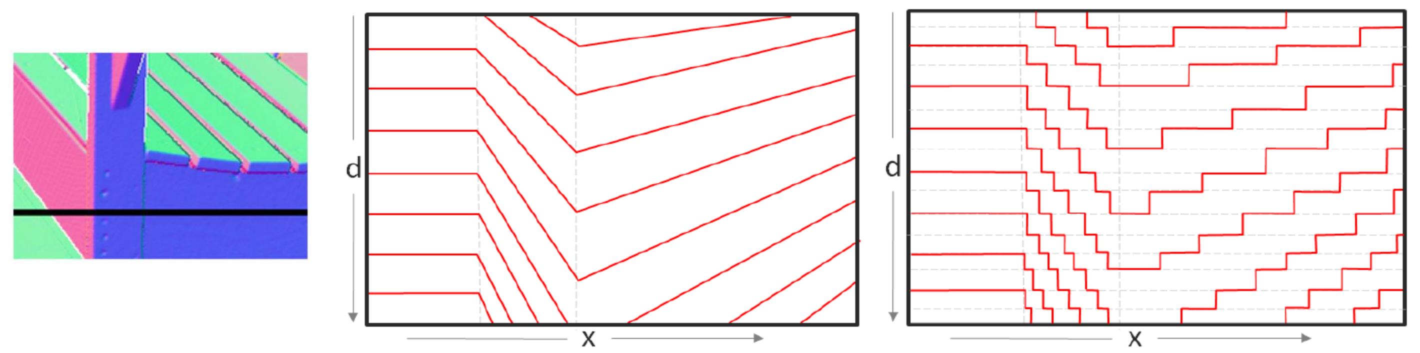

Experiments

- SGM vs SGM-P (various forms of prior)

2D Prior (offset image representation)
- SGM-EPi – Estimated segmented planes
- SGM-GS – GT surface
- SGM-GP – GT surface, planar approximation
- SGM-GNi – GT normals (fixed-z "strawman")

3D Prior (offset volume representation)
- SGM-EPv – Estimated overlapping planes
- SGM-GNv – GT normals (accurate version)
- SGM-MW – Manhattan-world prior

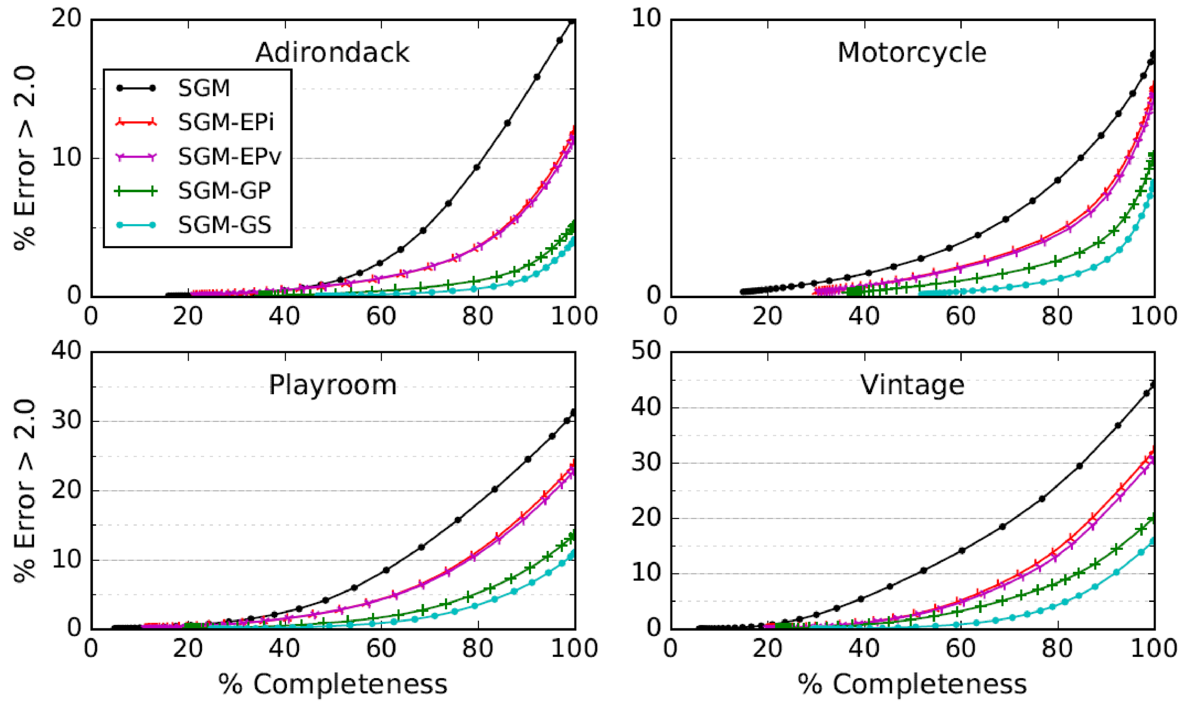

- Matching costs (NCC, MC-CNN)
- Improvements for slanted textureless surfaces; never hurts performance

Middlebury v3 Test Set Scores

Cost	SGM		SGM-EPi		SGM-EPv		SGM-GS	
	avg err	rank	avg err	rank	avg err	rank	avg err	rank
NCC	18.9	28	16.3	24	16.2	24	8.27	4
gain over SGM	-		14%		14%		56%	
MC-CNN	15.6	23	13.7	18	13.5	18	7.31	2
gain over SGM	-		12%		13%		53%	

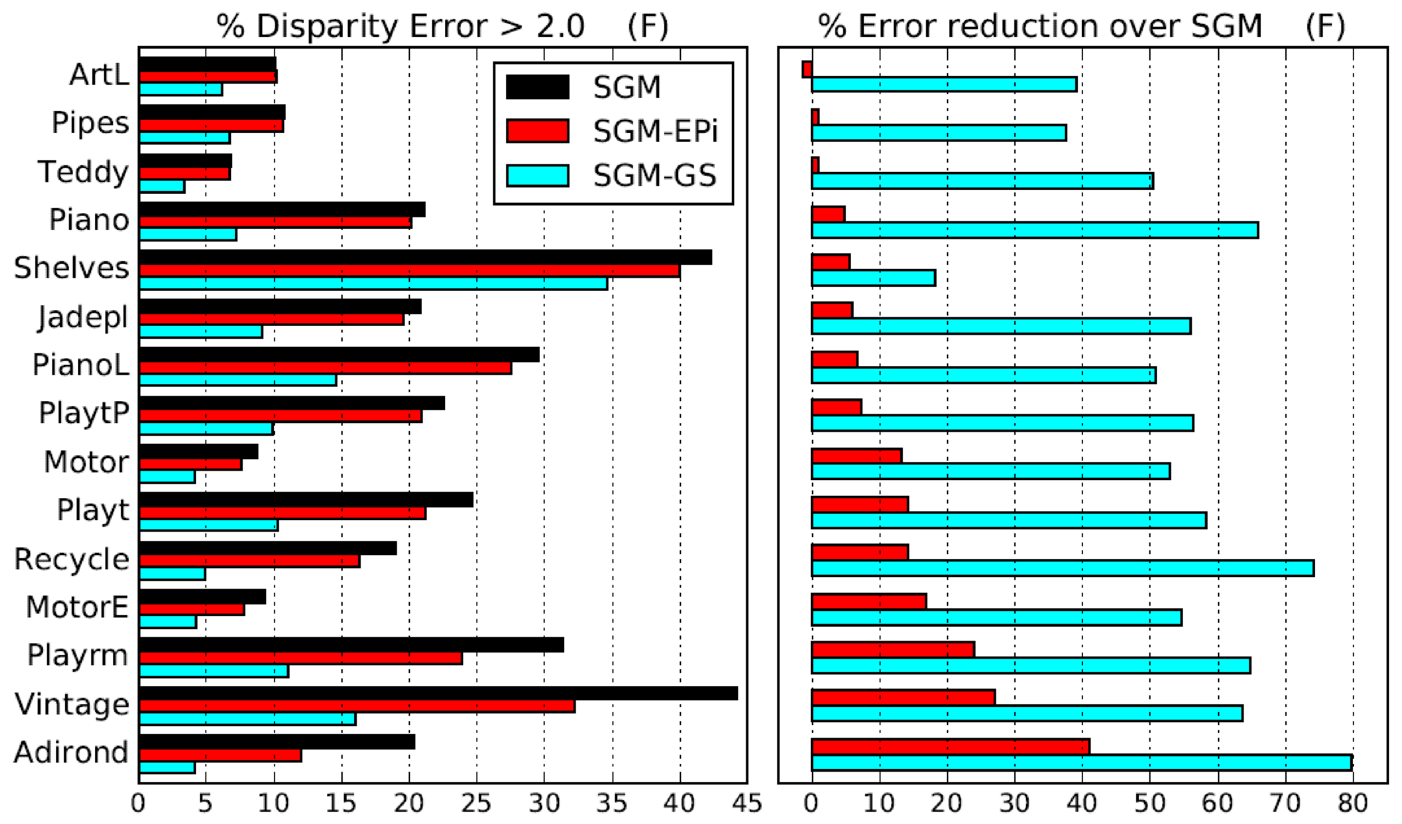

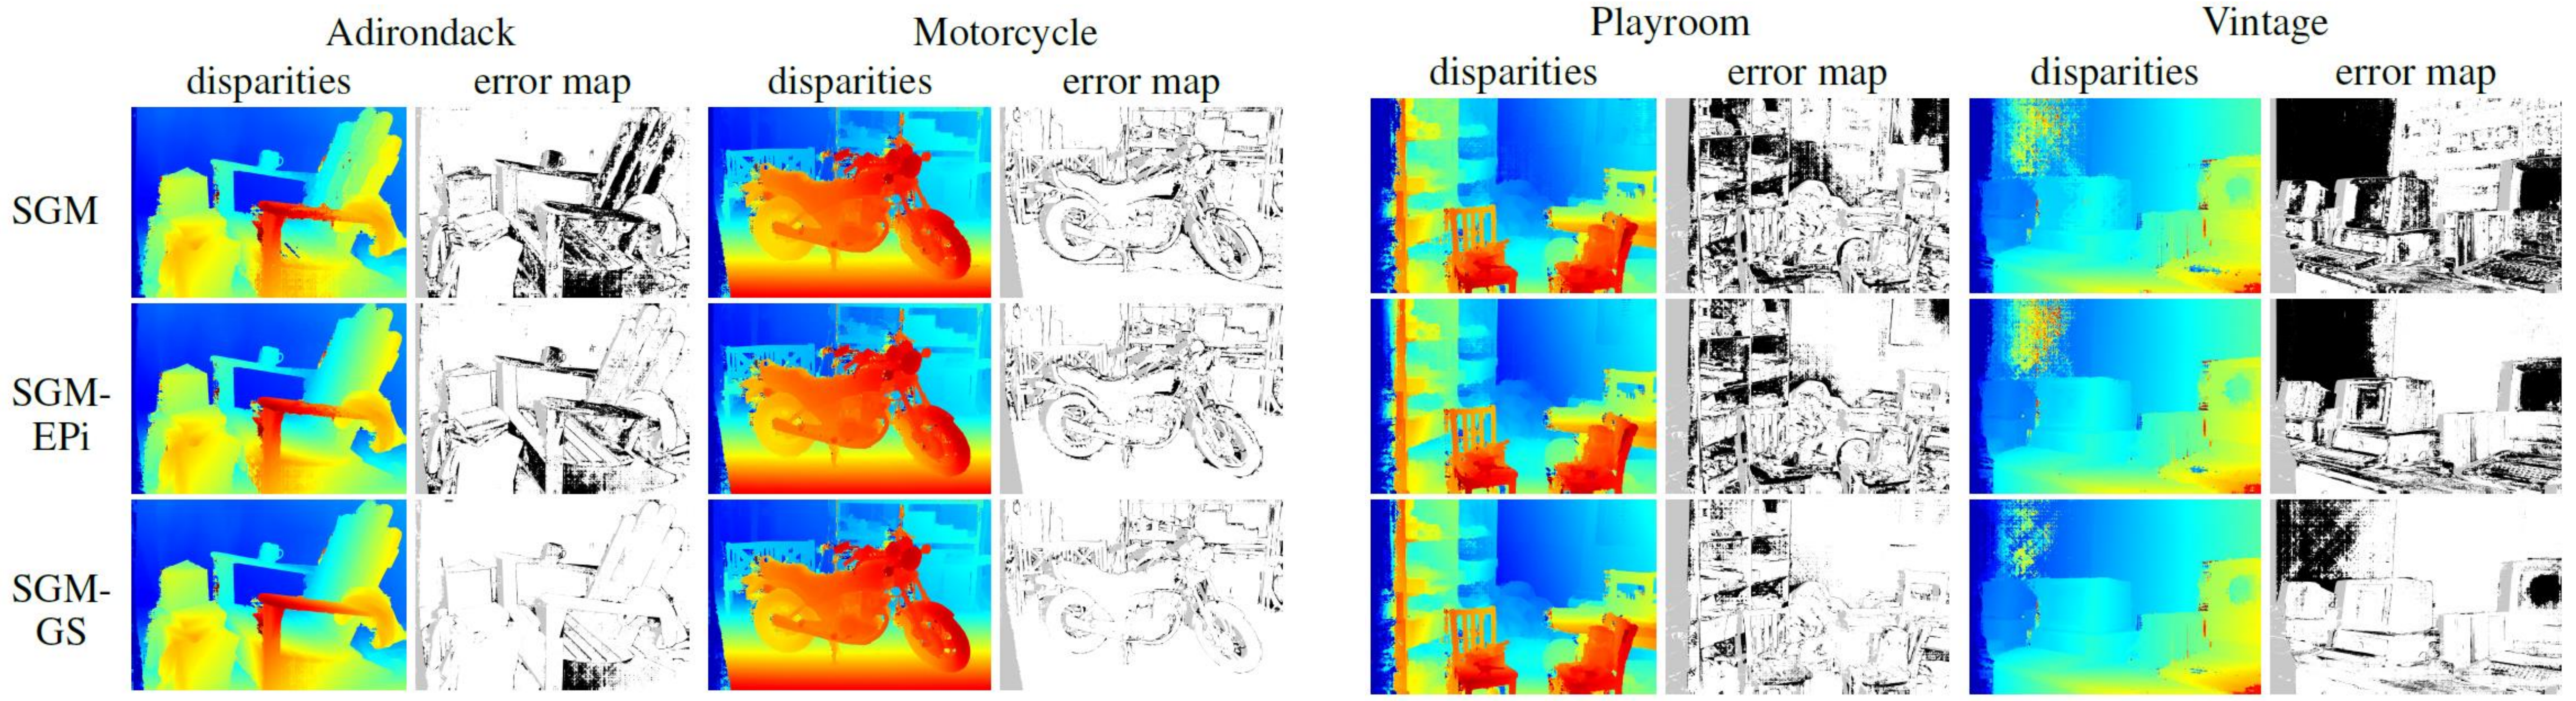

Summary

- Extension of SGM; uses precomputed surface orientation priors
- Huge performance gains for slanted untextured scenes
- Soft constraint, doesn't hurt performance; small overhead
- Explore predicting more accurate orientation priors in future.