Continuous Markov Random Fields for Robust Stereo Estimation

Koichiro Yamaguchi[1,2] Tamir Hazan[1] David McAllester[1] Raquel Urtasun[1]

[1] TTI-Chicago

[2] Toyota Central R&D Labs., Inc.

Goal

Dense stereo for high-resolution real-world images

Middlebury low-resolution dataset [Scharstein and Szeliski 2001]

- Low resolution
- Laboratory environment

State-of-the-art algorithms
Average error:2 – 3 %

KITTI Vision Benchmark Suite [Geiger, et al. 2012]

High-resolution realistic dataset

Difficulties

- Large number of disparity labels
- Textureless regions
- Strong slants

Related Work

Pixel-based MRF

Very local smoothness at pixel level

Slanted-plane MRF [Birchfield and Tomasi 1999]

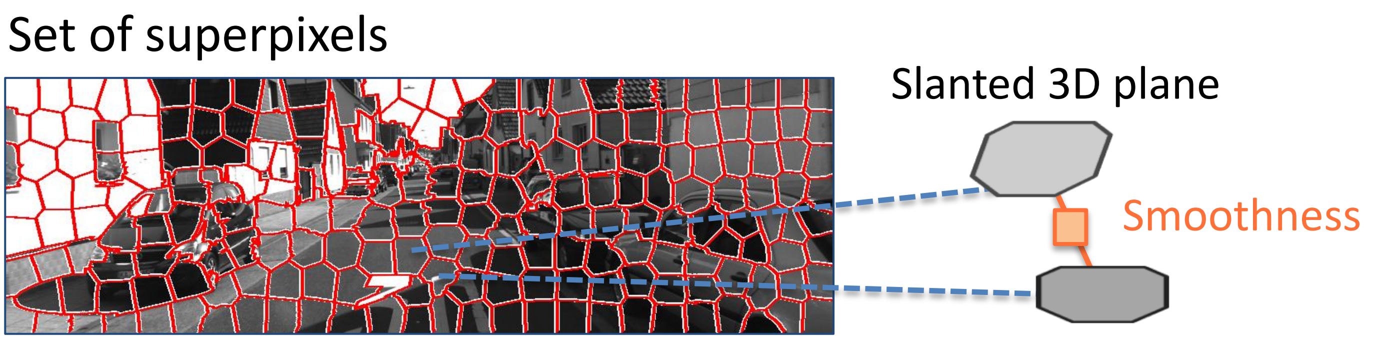

Continuous MRF is computationally challenging

[Bleyer, et al. 2010] 1 hour for low-resolution Middlebury image

Our Approach

Novel model for slanted-plane MRF

Introduce boundary labels and junction feasibility

Inference using Particle Convex Belief Propagation

Perform with reasonable running time

Our model

Random Variables

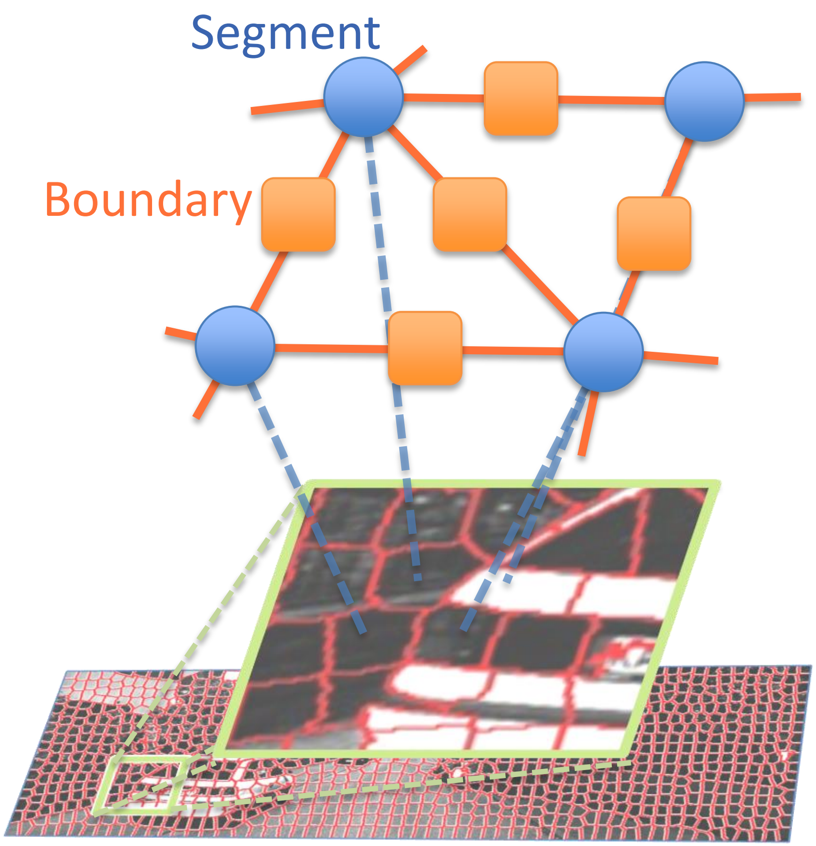

Superpixels (UCM [Arbelaez, et al. 2011] and SLIC [Achanta, et al. 2010])

Segment variable $\mathbf{y}_i = (\alpha_i, \beta_i, \gamma_i)$

Slanted 3D plane of segment

Continuous variable

Boundary variable o_{ij}

Relationship between segments

4 states

Occlusion Hinge Coplanar

Discrete variable

Energy Function

$$E(\mathbf{y}, \mathbf{o}) = E_{\text{color}}(\mathbf{o}) + E_{\text{match}}(\mathbf{y}, \mathbf{o}) + E_{\text{compatibility}}(\mathbf{y}, \mathbf{o}) + E_{\text{junction}}(\mathbf{o})$$

$\mathbf{y}$: set of all 3D slanted planes

$\mathbf{o}$: set of all boundary variables

Energy Function

$$E(\mathbf{y},\mathbf{o}) = E_{\text{color}}(\mathbf{o}) + E_{\text{match}}(\mathbf{y},\mathbf{o}) + E_{\text{compatibility}}(\mathbf{y},\mathbf{o}) + E_{\text{junction}}(\mathbf{o})$$

Color similarity energy

Similar color → Likely to be coplanar

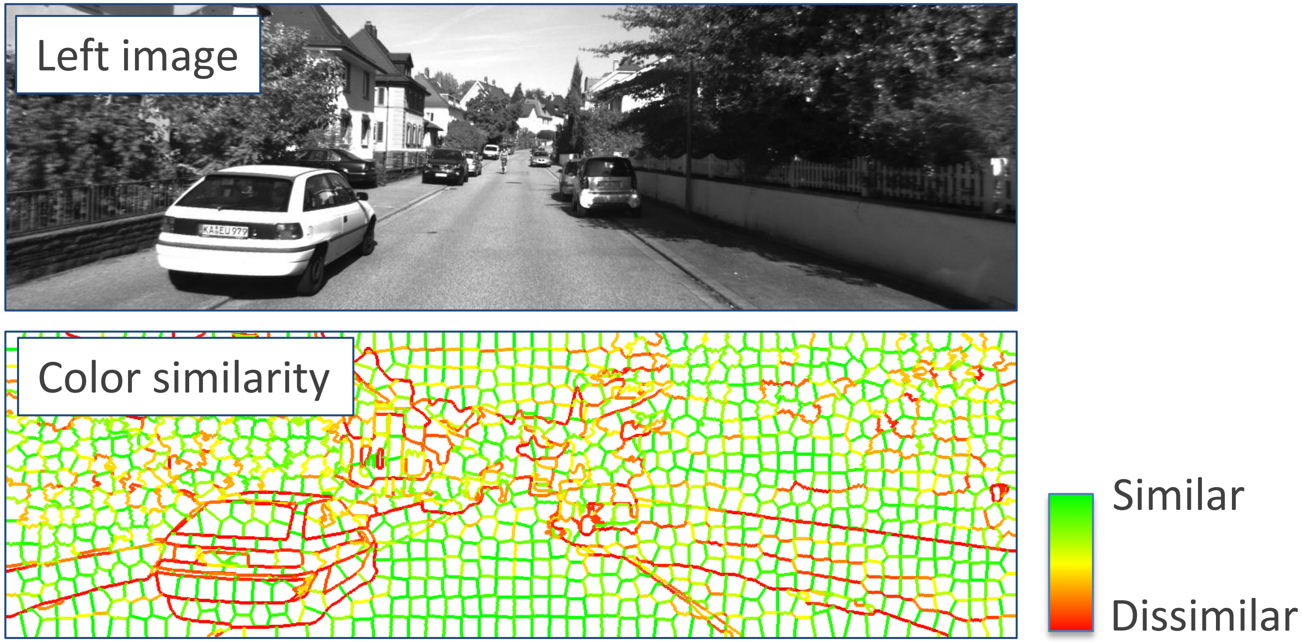

Energy Function

$$E(\mathbf{y}, \mathbf{o}) = E_{\text{color}}(\mathbf{o}) + E_{\text{match}}(\mathbf{y}, \mathbf{o}) + E_{\text{compatibility}}(\mathbf{y}, \mathbf{o}) + E_{\text{junction}}(\mathbf{o})$$

Matching energy

Agreement with result of input disparity map

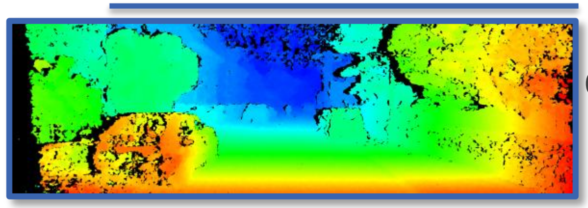

Computed by any matching method
(Modified semi-global matching)

Truncated quadratic function $\phi_i^{\text{TP}}(\mathbf{p}, \mathbf{y}_i, K) = \min\left(|\mathcal{D}(\mathbf{p}) - \hat{d}_i(\mathbf{p}, \mathbf{y}_i)|, K\right)^2$

Disparity map　Slanted plane

On boundary

"Occlusion" – Foreground segment owns boundary

Energy Function

$$E(\mathbf{y},\mathbf{o}) = E_{\text{color}}(\mathbf{o}) + E_{\text{match}}(\mathbf{y},\mathbf{o}) + E_{\text{compatibility}}(\mathbf{y},\mathbf{o}) + E_{\text{junction}}(\mathbf{o})$$

Compatibility energy

(1) Preference of boundary label (Coplanar > Hinge > Occlusion)

Impose penalty $\lambda_{\text{occ}} > \lambda_{\text{hinge}} > 0$

(2) Boundary labels ⟷ (match) ⟷ Slanted planes

- "Occlusion" ⟷ $\hat{d}_{\text{front}}(\mathbf{p}) > \hat{d}_{\text{back}}(\mathbf{p})$
- "Hinge" ⟷ $\hat{d}_i(\mathbf{p}) = \hat{d}_j(\mathbf{p})$ on boundary
- "Coplanar" ⟷ $\hat{d}_i(\mathbf{p}) = \hat{d}_j(\mathbf{p})$ in both segments

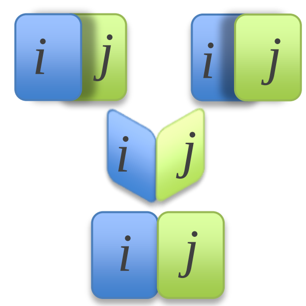

Energy Function

$$E(\mathbf{y},\mathbf{o}) = E_{\text{color}}(\mathbf{o}) + E_{\text{match}}(\mathbf{y},\mathbf{o}) + E_{\text{compatibility}}(\mathbf{y},\mathbf{o}) + E_{\text{junction}}(\mathbf{o})$$

Junction energy

Occlusion boundary reasoning [Malik 1987]

Penalize impossible junctions

Impossible cases

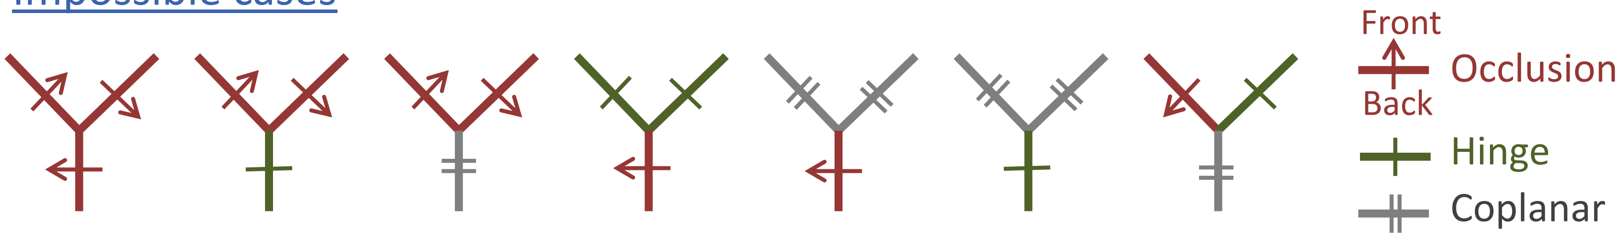

Hybrid MRF

defined over continuous variables $\mathbf{y}$ and discrete variables $\mathbf{o}$

Inference / Learning

Inference

Particle Convex Belief Propagation (PCBP)

3D slanted planes $\mathbf{y}$ → Discretized as particles

Algorithm

- Initialize slanted planes $\mathbf{y}_i^0 = (\alpha_i^0, \beta_i^0, \gamma_i^0)$ via local fitting
- **for** $t = 1$ to T
  - Sample $\alpha_i, \beta_i, \gamma_i$
  - $(\mathbf{o}^t, \mathbf{y}^t)$ ← Solve discretized MRF using convex BP [Schwing, et al. 2011]
- Return $(\mathbf{o}^T, \mathbf{y}^T)$

Learning

Discretized $\mathbf{y}$ as particles

Use training algorithm based on primal-dual approximate inference [Hazan and Urtasun 2010]

Experiments

Middlebury high-resolution images [Scharstein and Pal 2007]

- Laboratory environment
- High-resolution (1239x1038 pixels)
- 5 train / 9 test images

KITTI dataset [Geiger, et al. 2012]

- Real-world stereo dataset
- Accurate ground truth
- High-resolution (1237x374 pixels)
- 10 train / 174 validation / 195 test images

Evaluation - Middlebury

Comparison on Middlebury high-resolution dataset

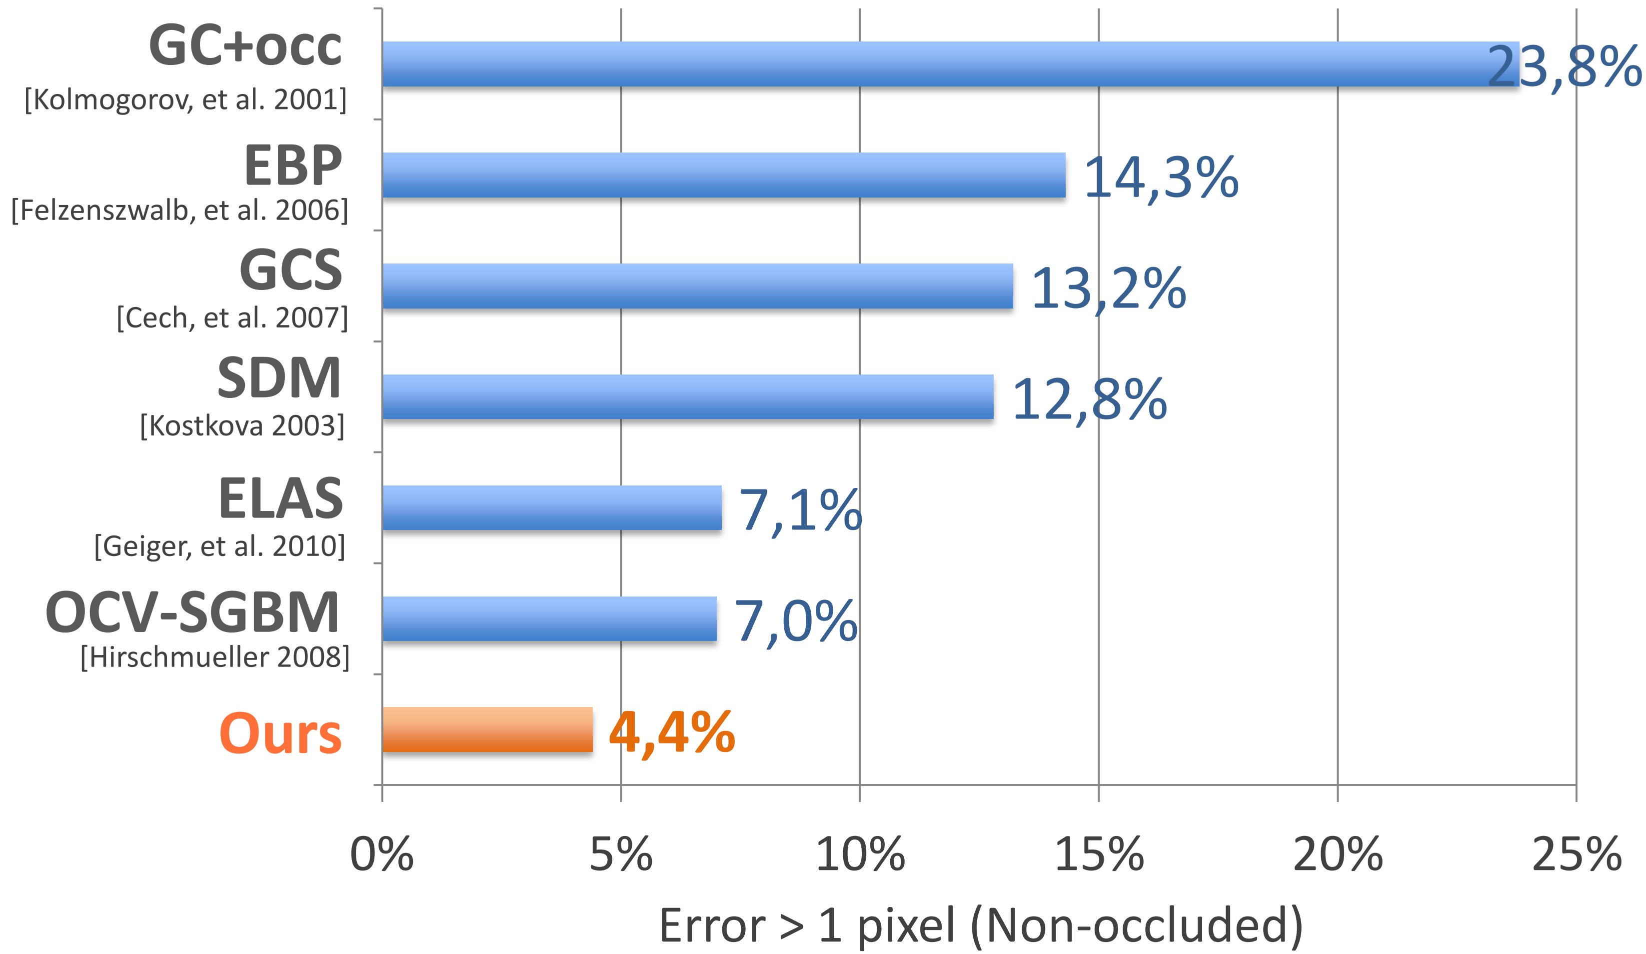

Result Example - Middlebury

Our inference

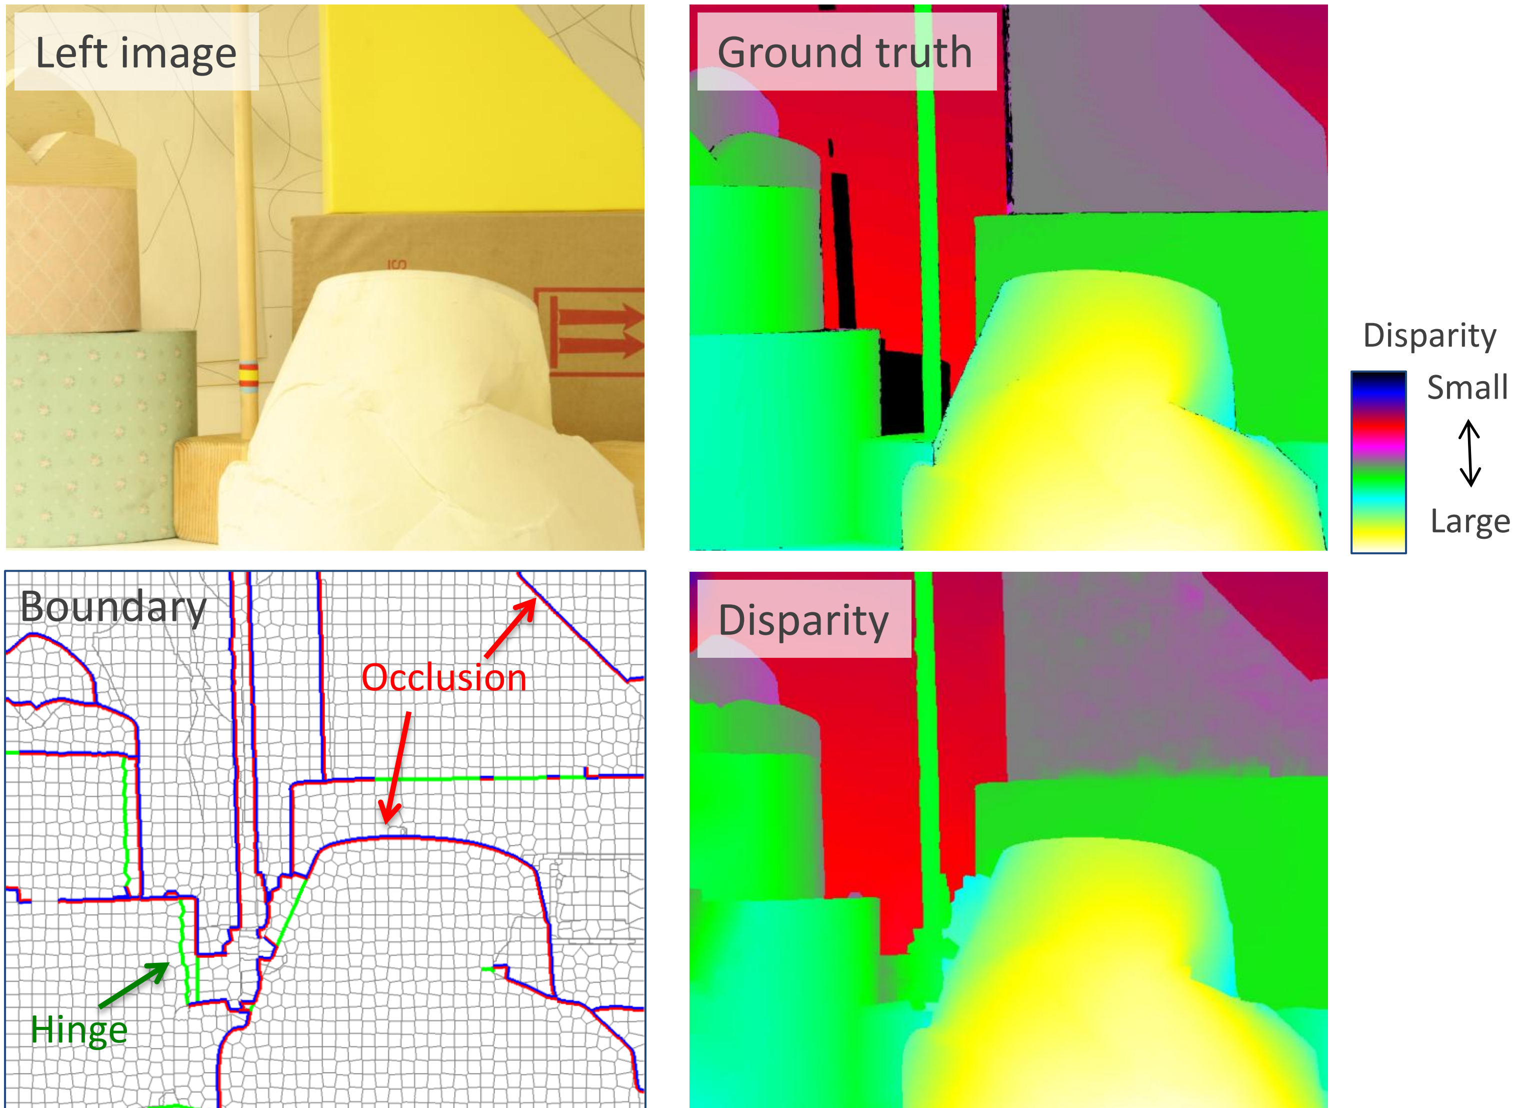

Result Example - Middlebury

OCV-SGBM

[Hirschmueller 2008]

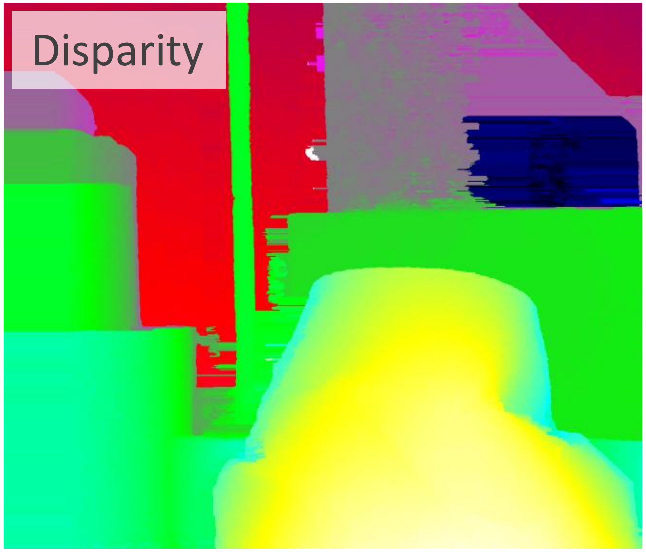

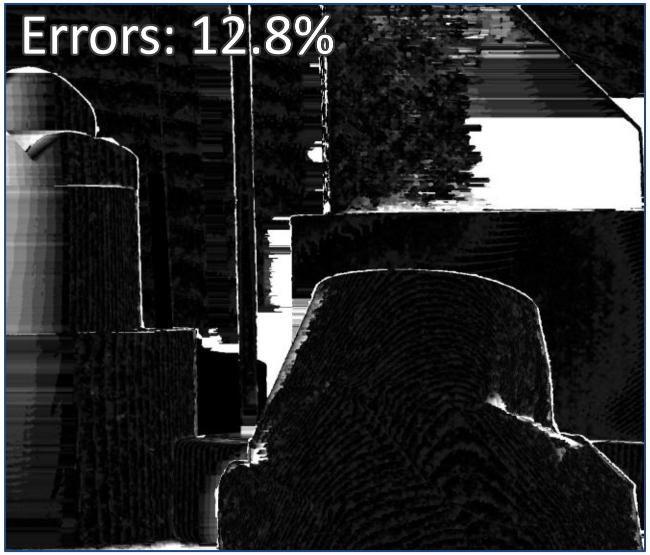

Ours

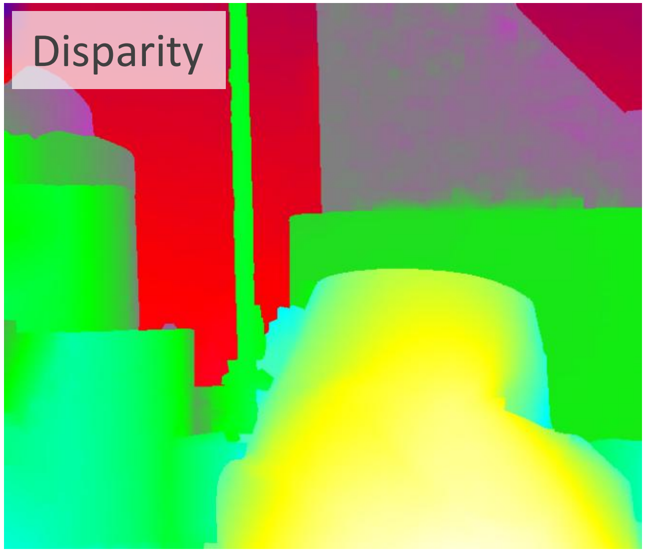

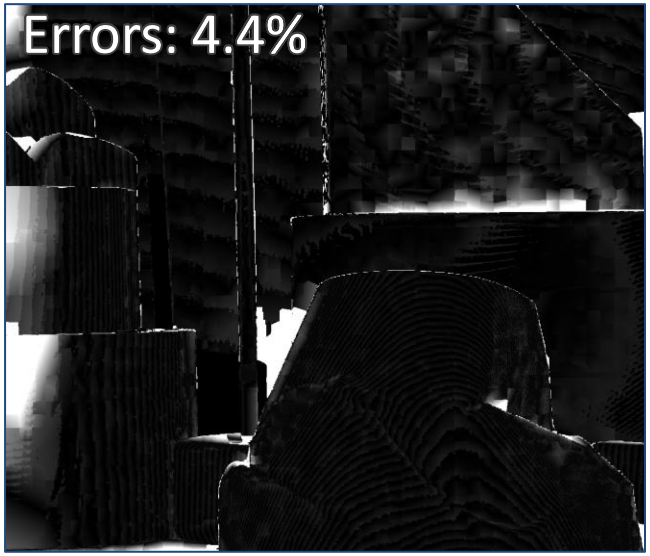

Evaluation - KITTI

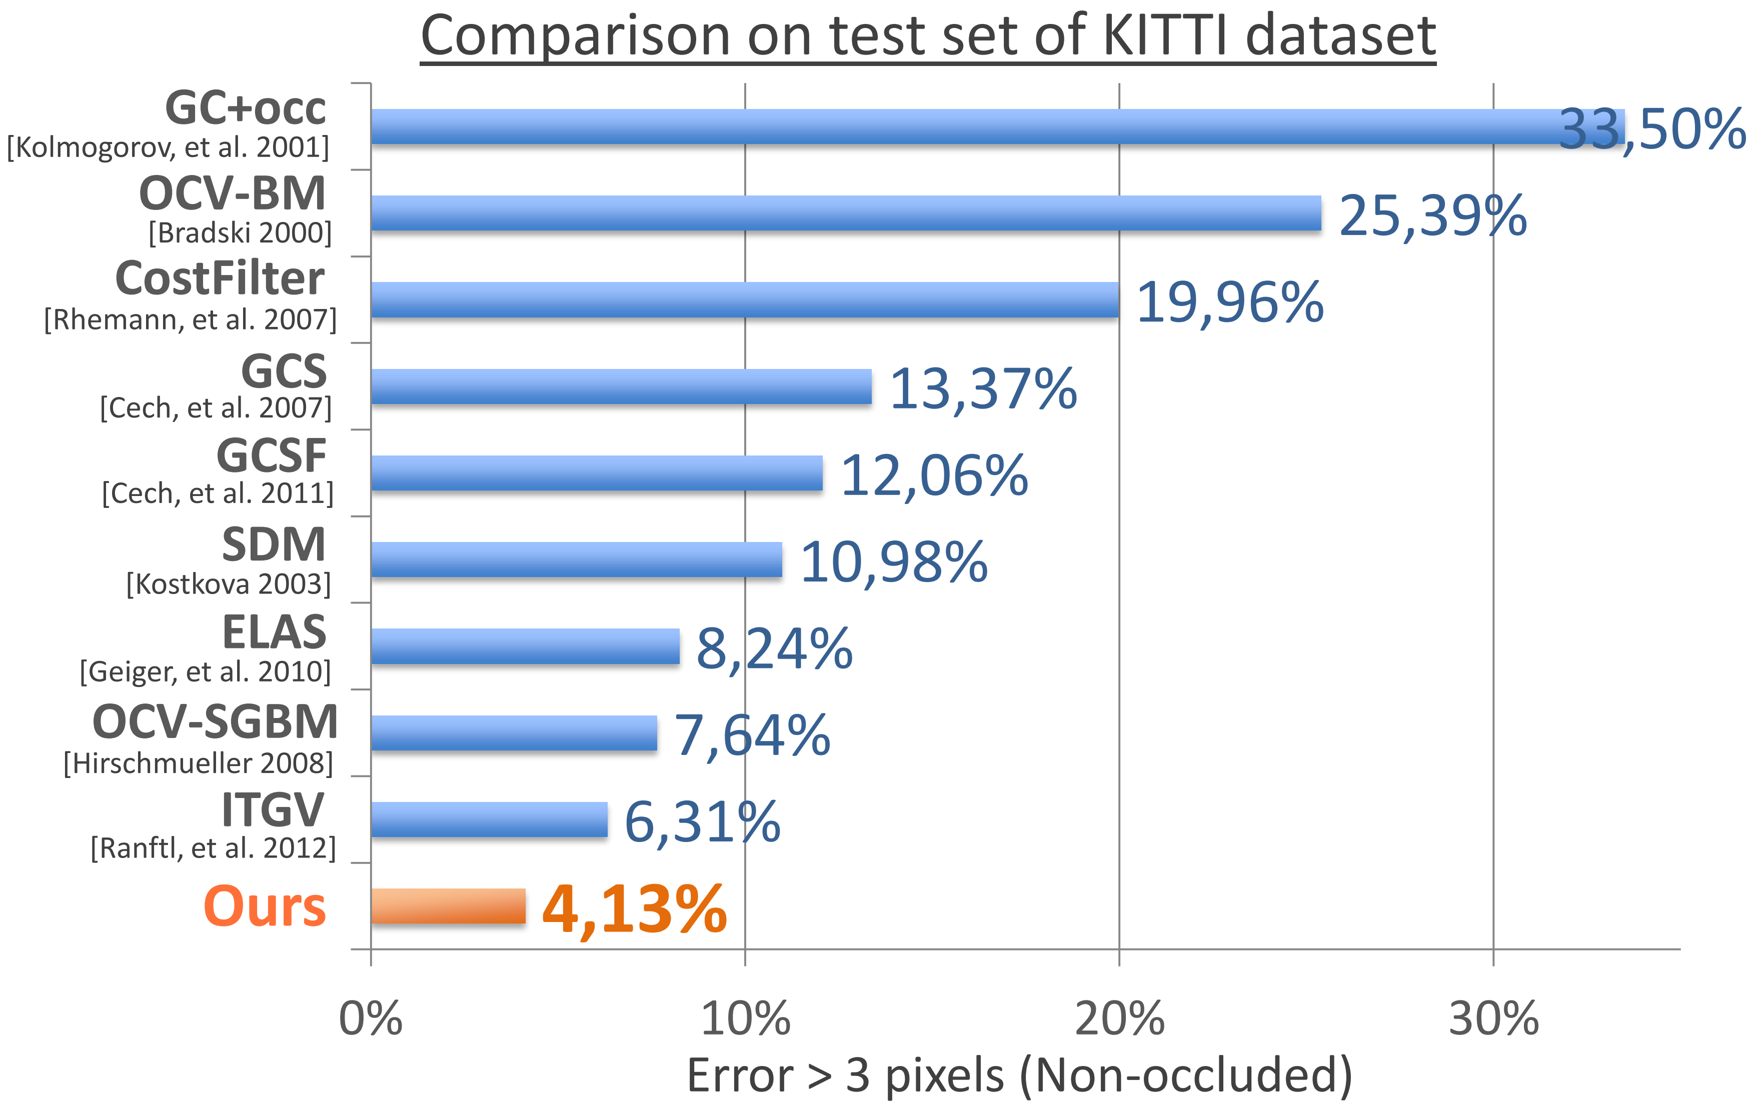

Result Examples - KITTI

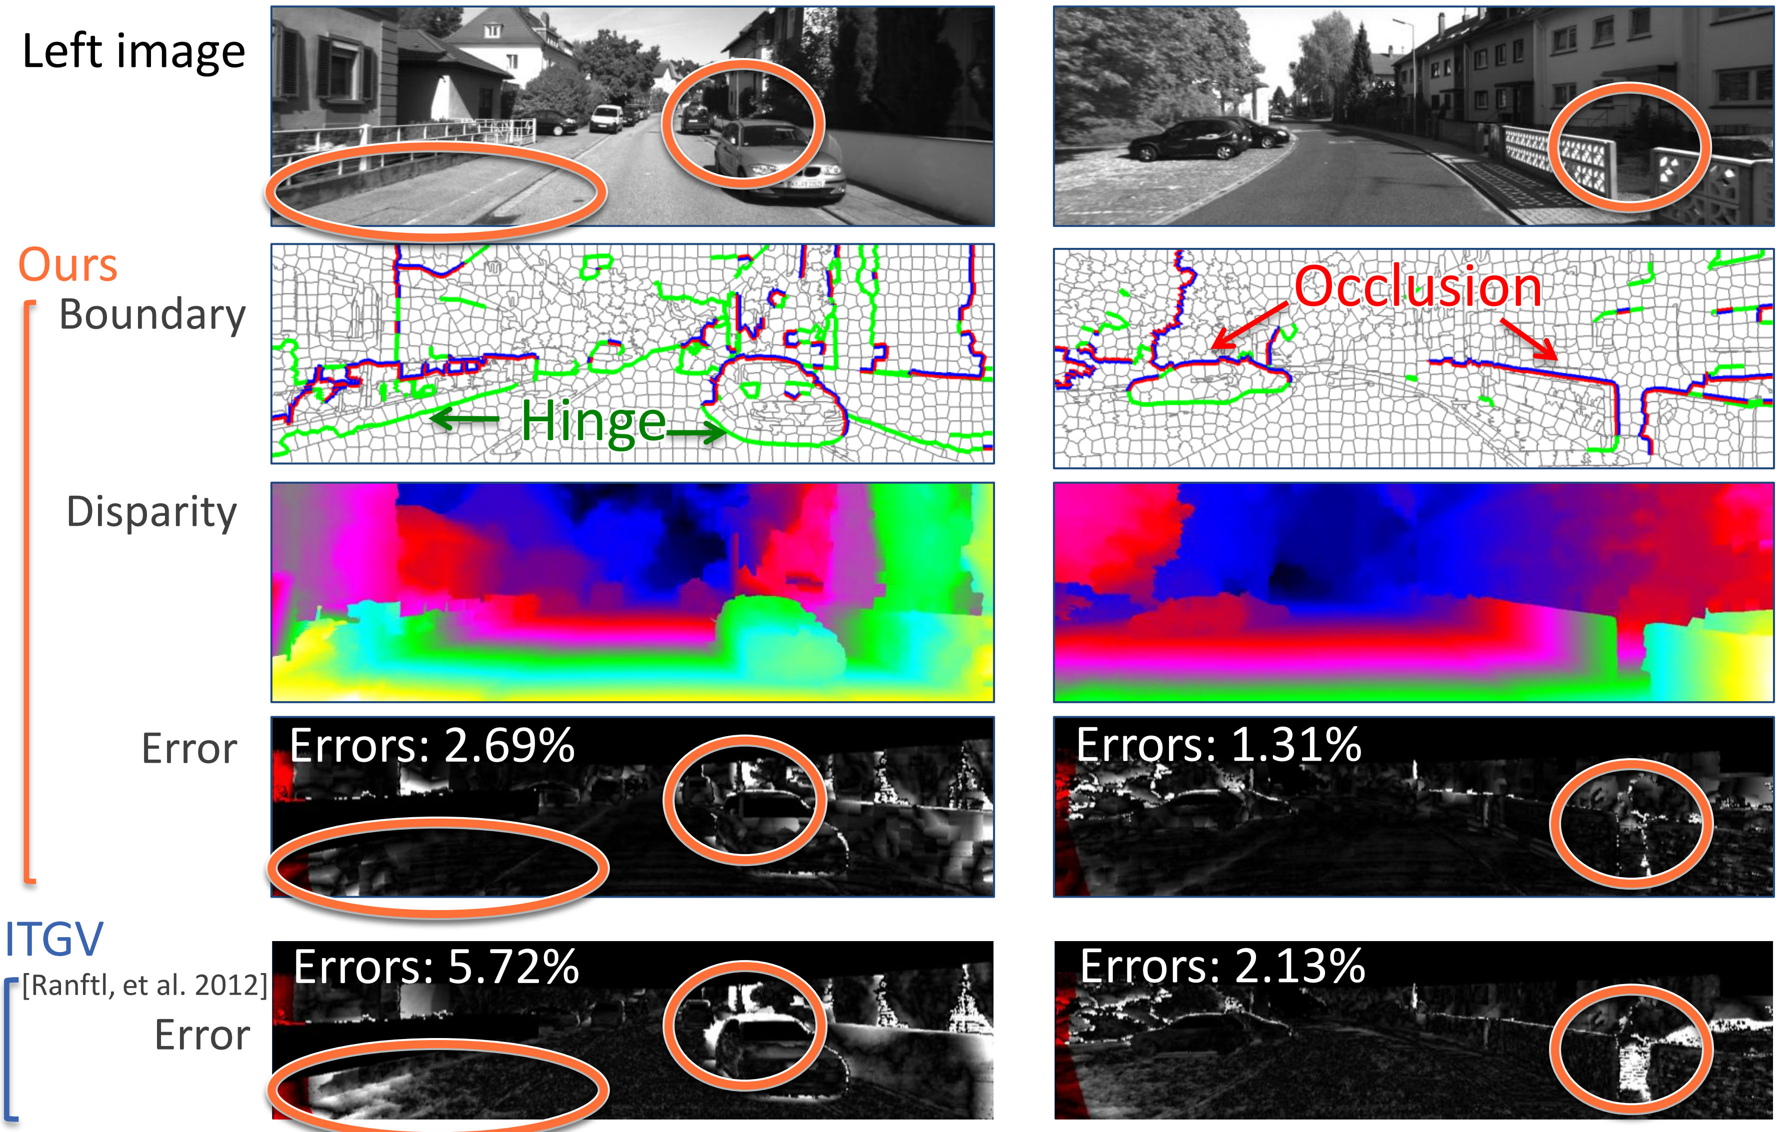

Training Set Size

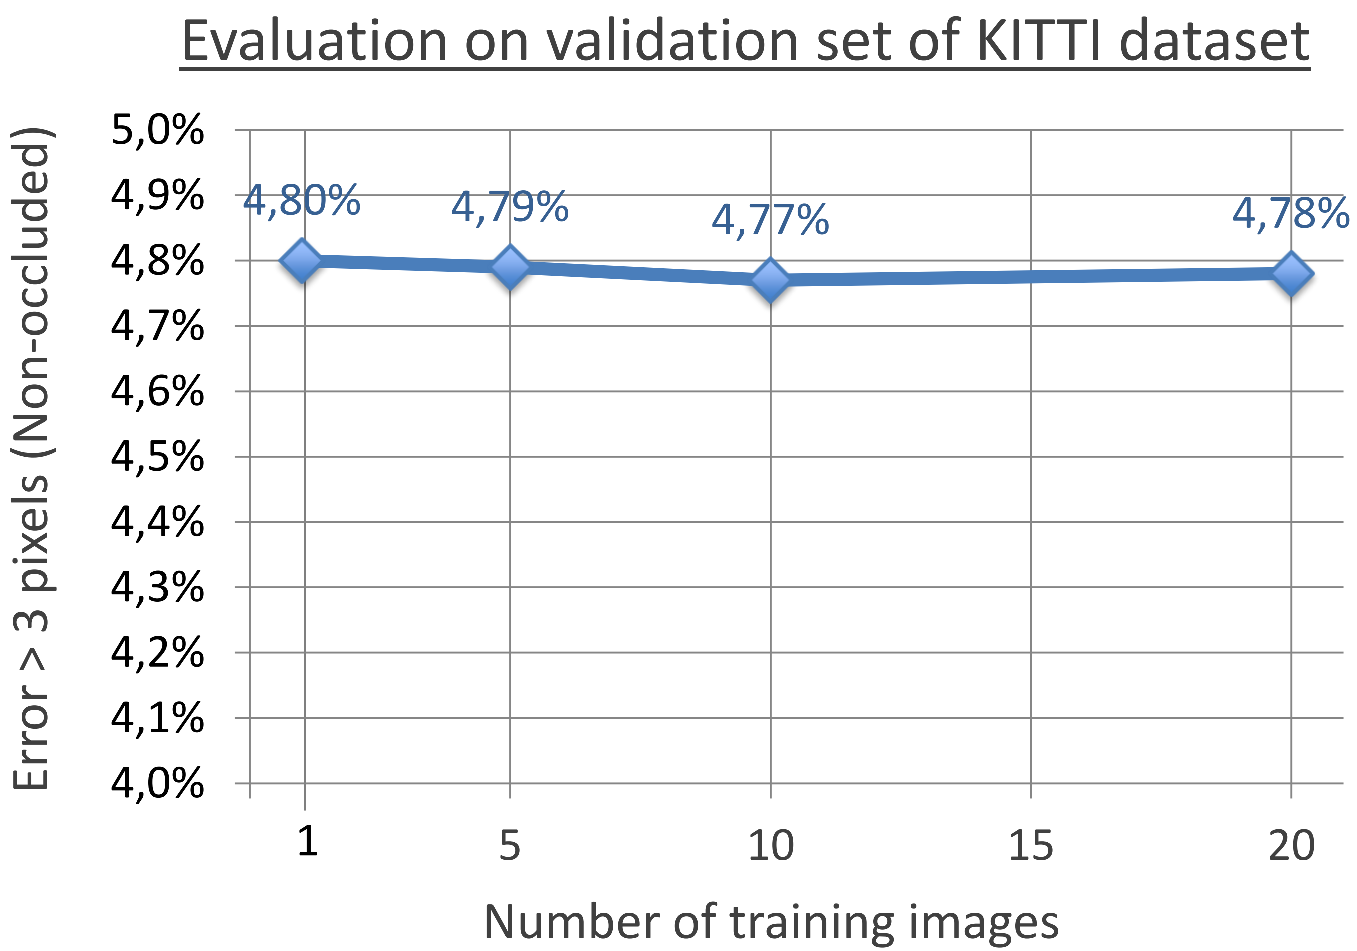

Importance of Energy Terms

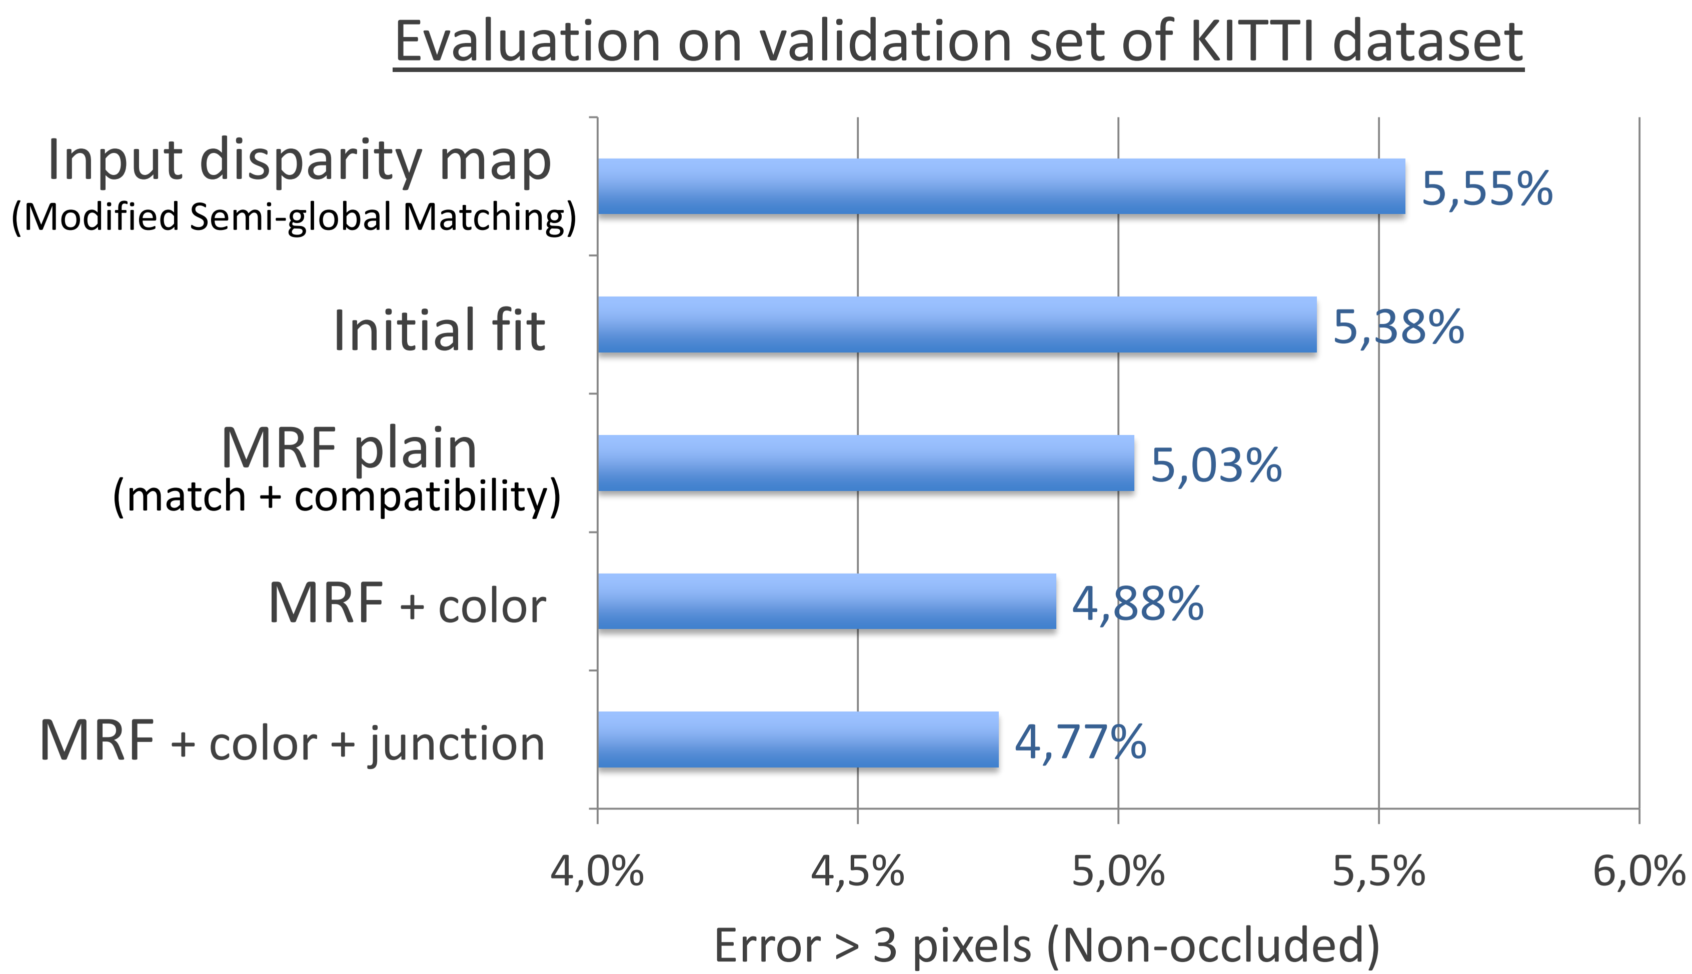

Conclusion

Novel slanted-plane MRF model

Estimate jointly occlusion boundaries and depth

Inference in hybrid MRF

Use particle convex belief propagation

Experiments on high resolution imagery

Outperform existing method

Future work

- Improve superpixel segmentation
- Investigate other potentials