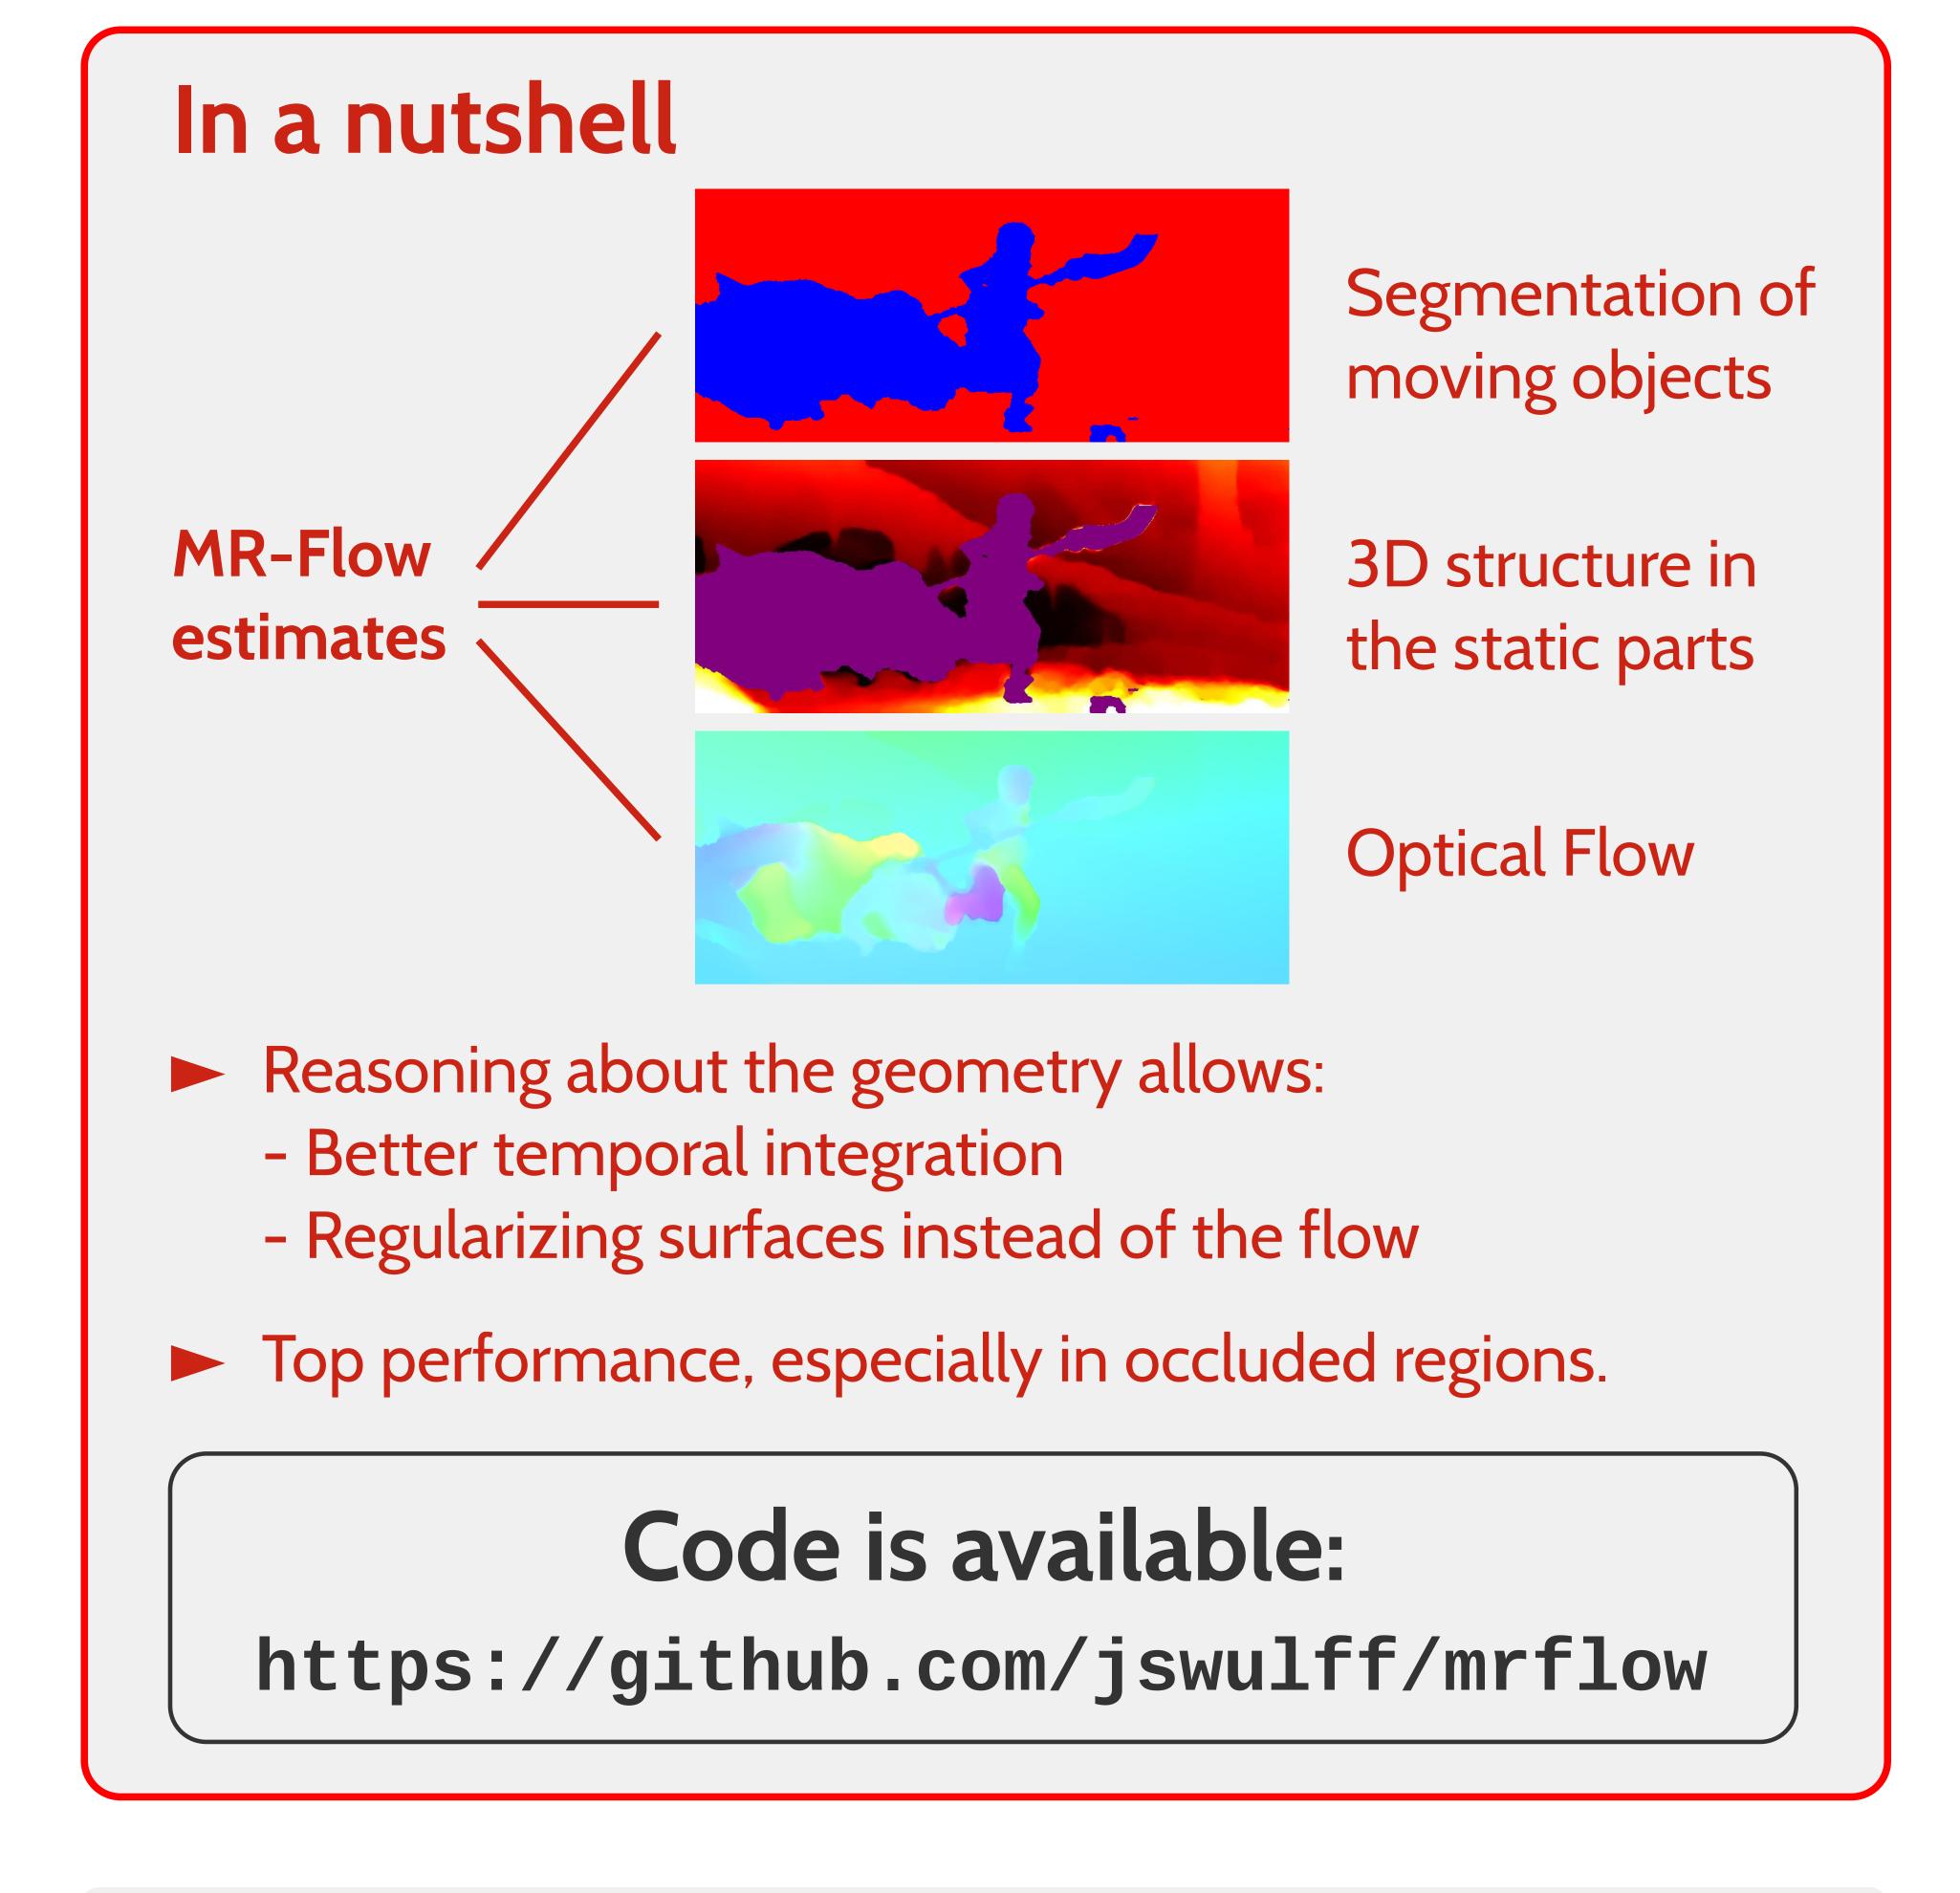


Optical Flow in Mostly Rigid Scenes

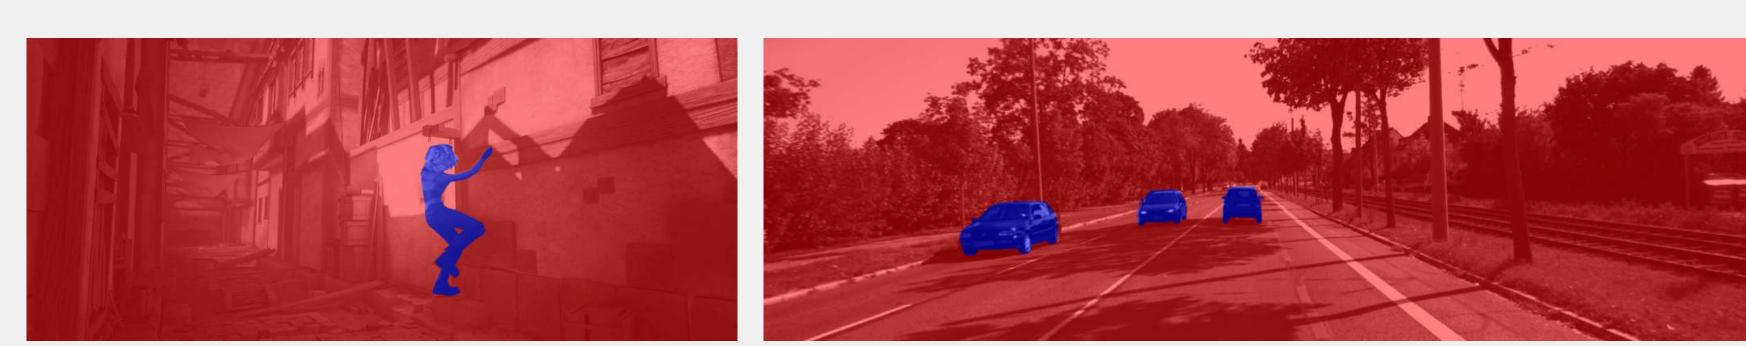
Jonas Wulff, Laura Sevilla-Lara, Michael J. Black

Max Planck Institute for Intelligent Systems



Introduction

Natural scenes are mostly static and do not move.



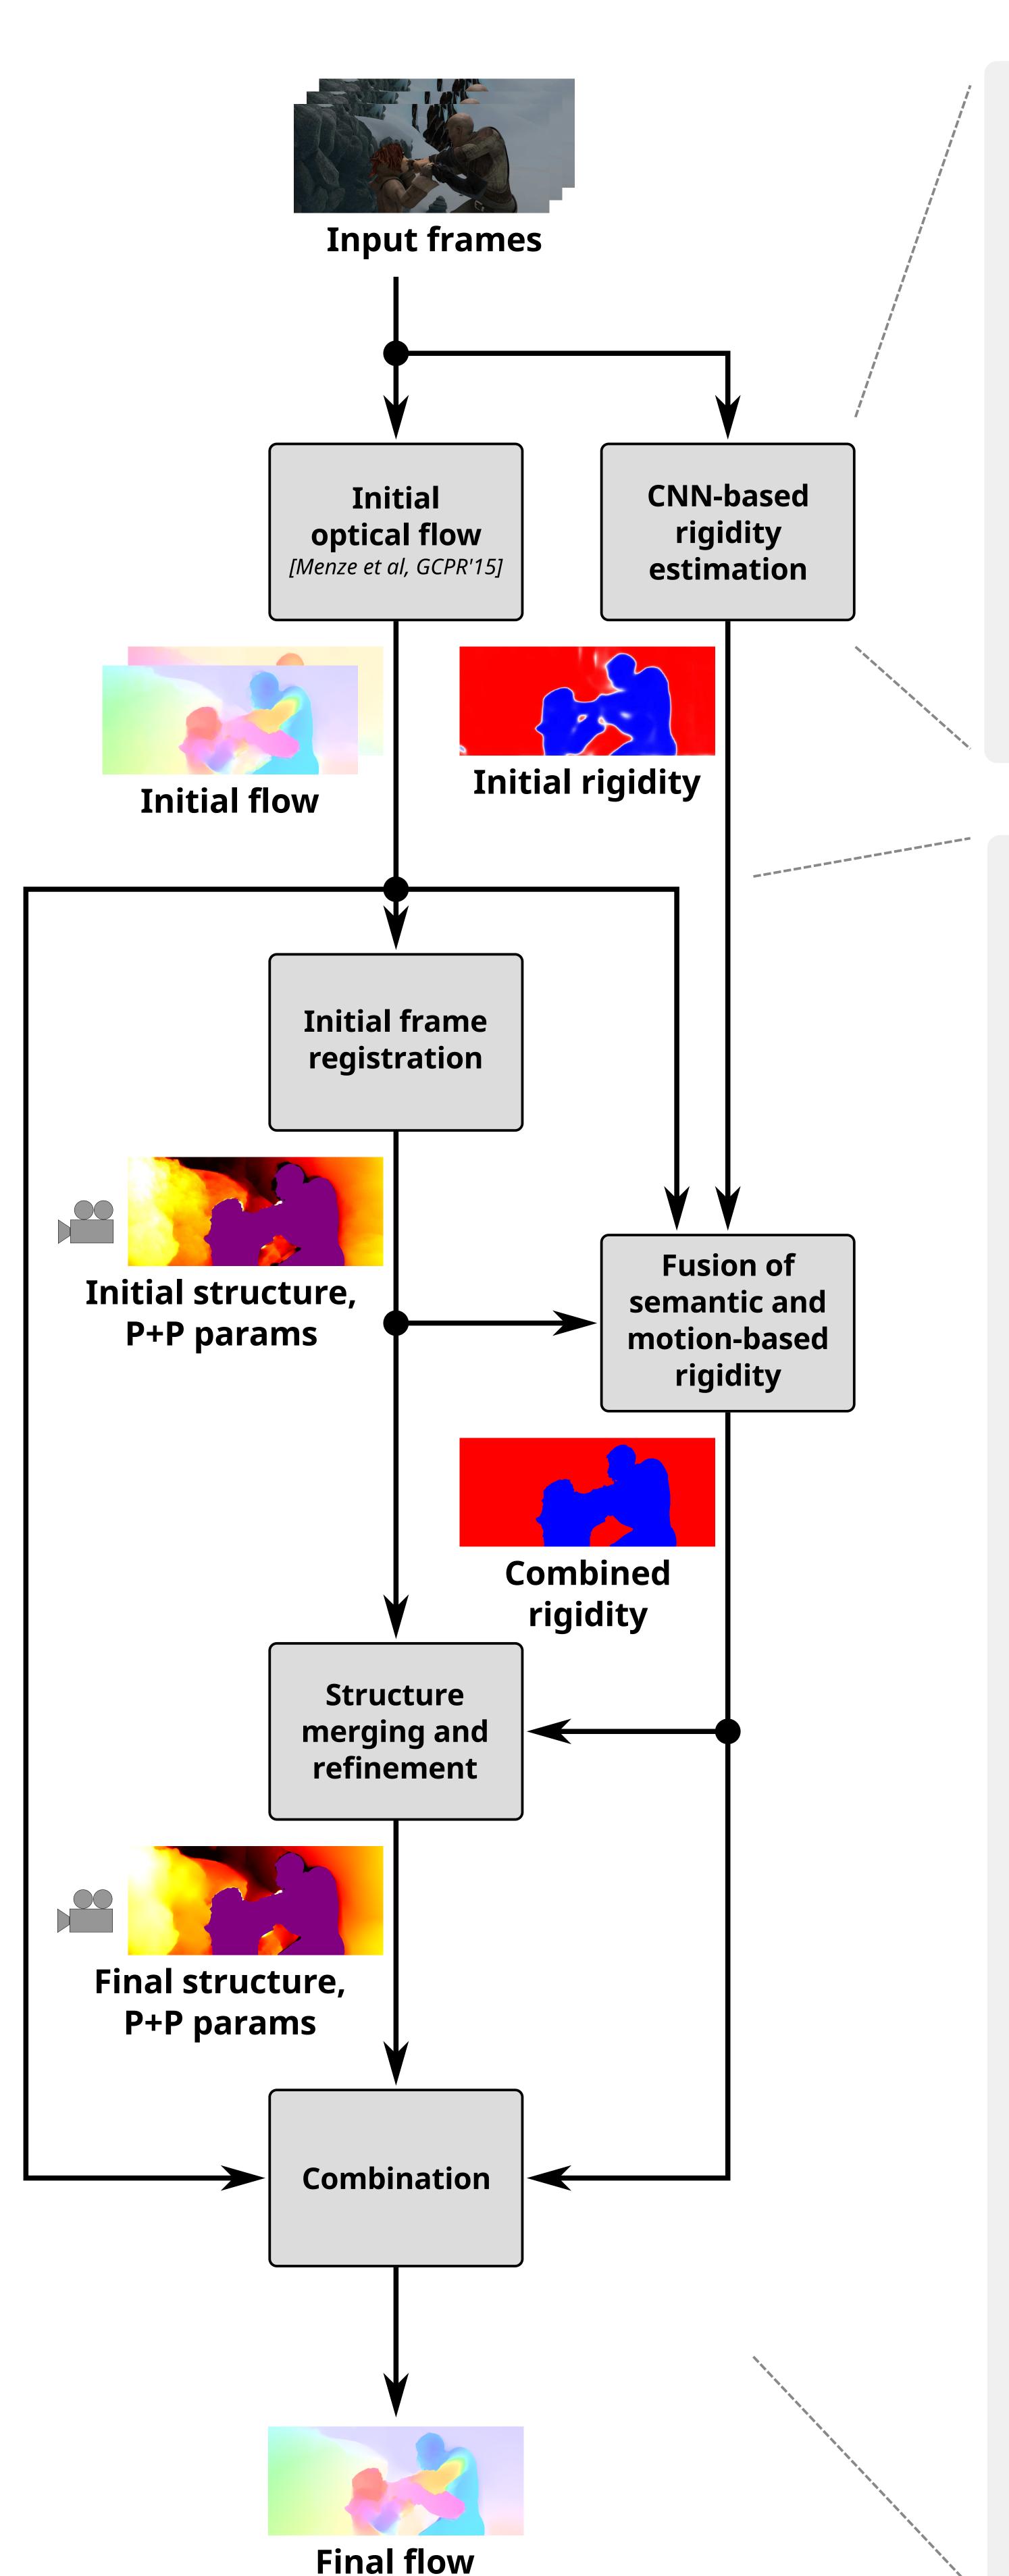
Sintel: 28% non-static

KITTI: 15 % non-static

- The optical flow in static parts is determined by the depth and camera motion.
- In static parts, the depth is constant in time.
- Moving surfaces correspond to objects, and recent progress in semantic segmentation yield good object segmentation results.

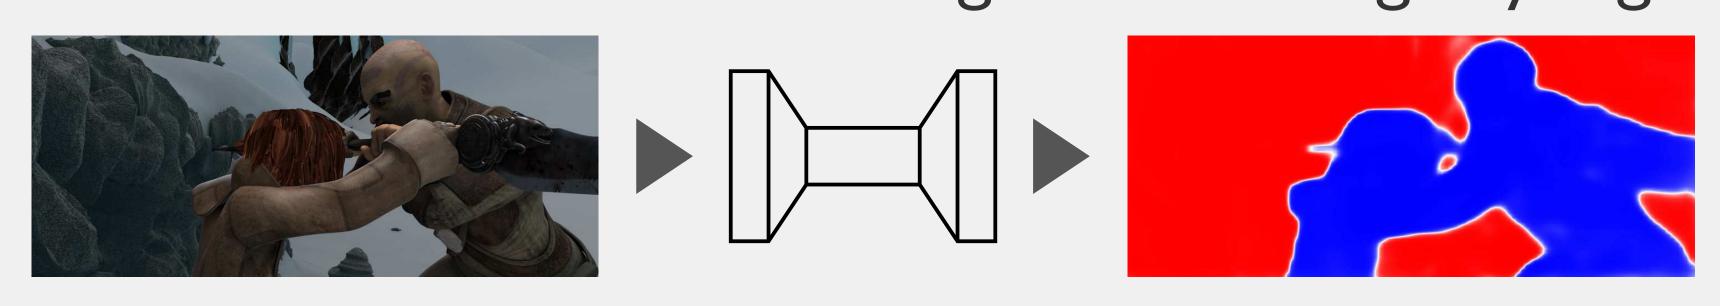
Basic idea:

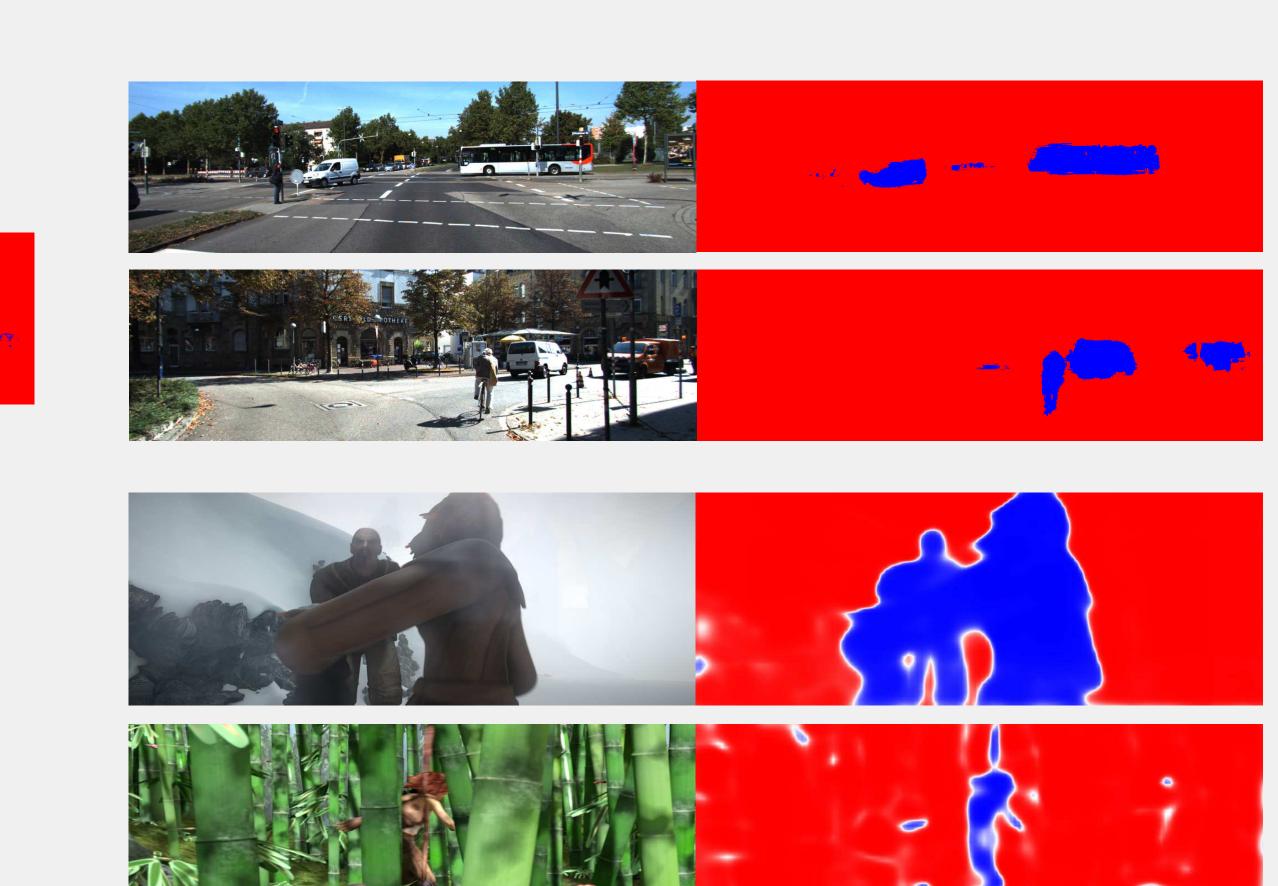
We split the scene into the static background and moving objects, and impose strong constraints wherever appropriate.



Semantic rigidity estimation

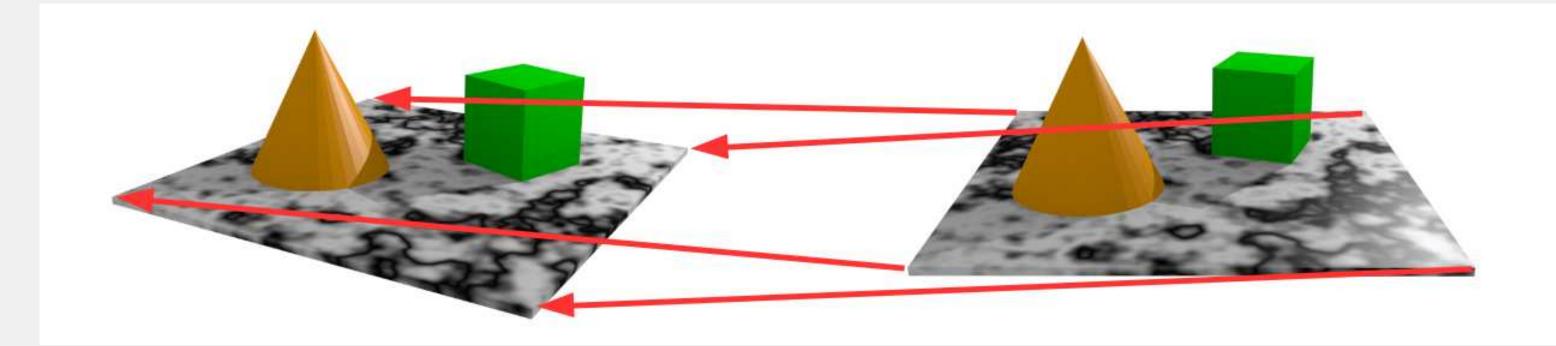
- ► KITTI: Segment using DeepLab [Chen et al, CoRR'14], assign rigidity to classes.
- MPI-Sintel: Pre-trained semantic segmentation does not work. Instead we re-train the CNN on ground truth rigidity segmentation.





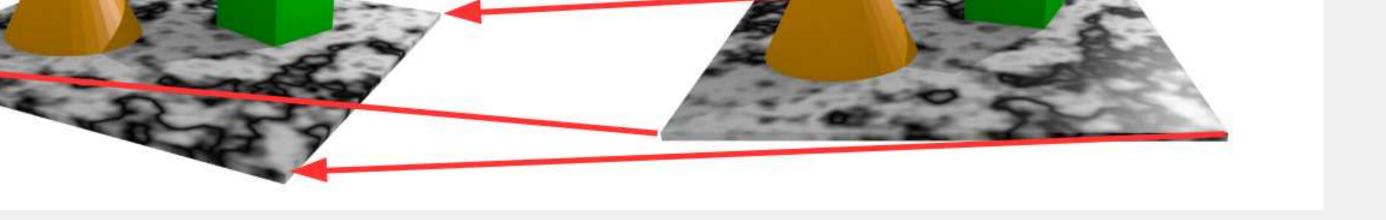
Motion estimation in the static scene

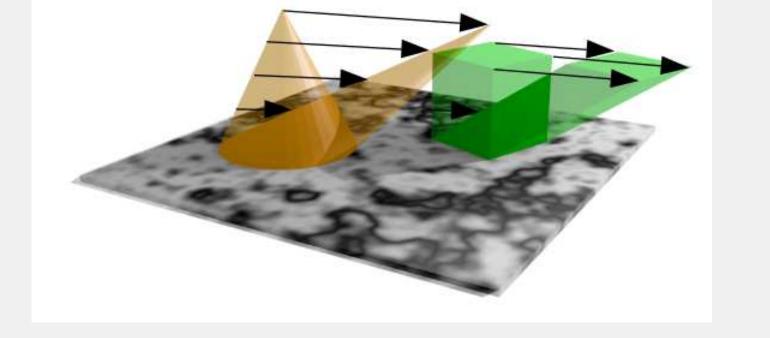
- SfM requires many frames, wide baselines <-> Optical Flow: only few frames and small baselines are available.
- Instead of SfM, we use Plane+Parallax [Sawhney, ECCV'94; Irani et al, ECCV'98].

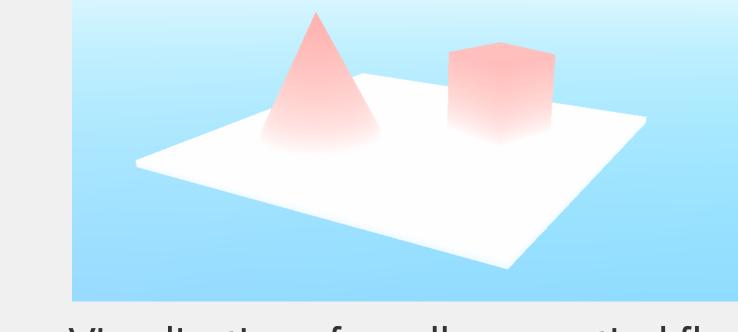


(1) Compute a homography to register frames to a common plane.

Definition of Parallax:







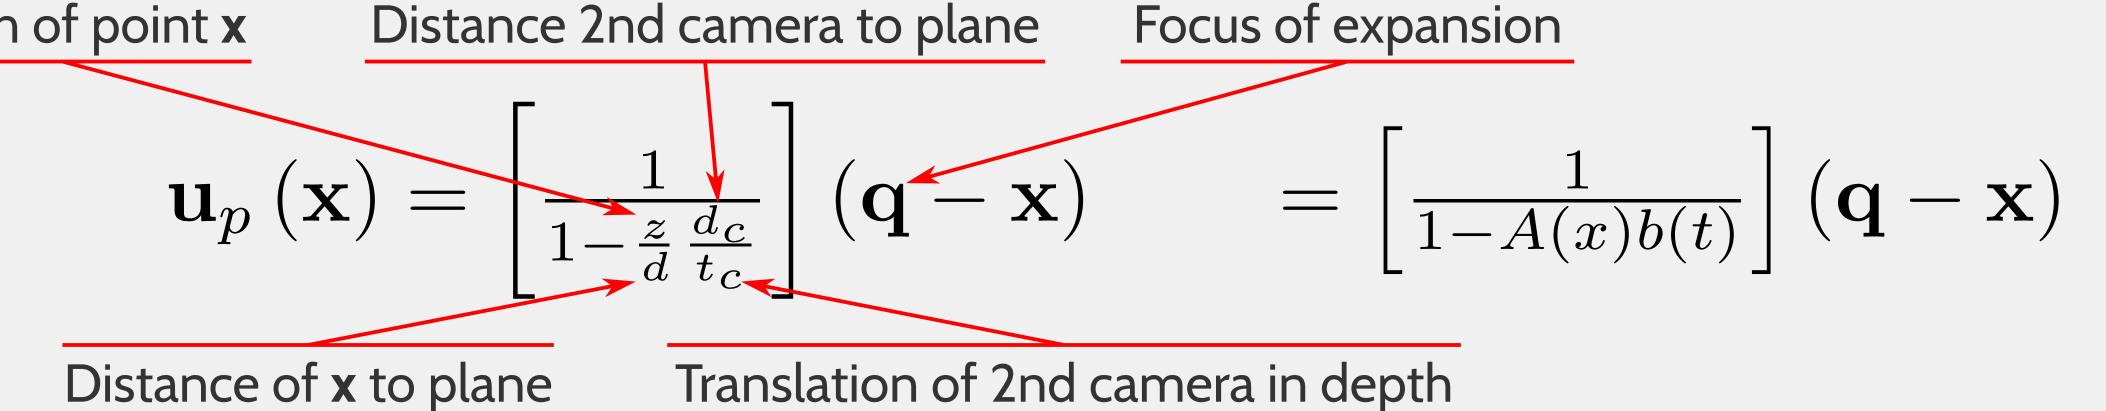
(2) Residual flow (=parallax) points to Epipole.

Visualization of parallax as optical flow

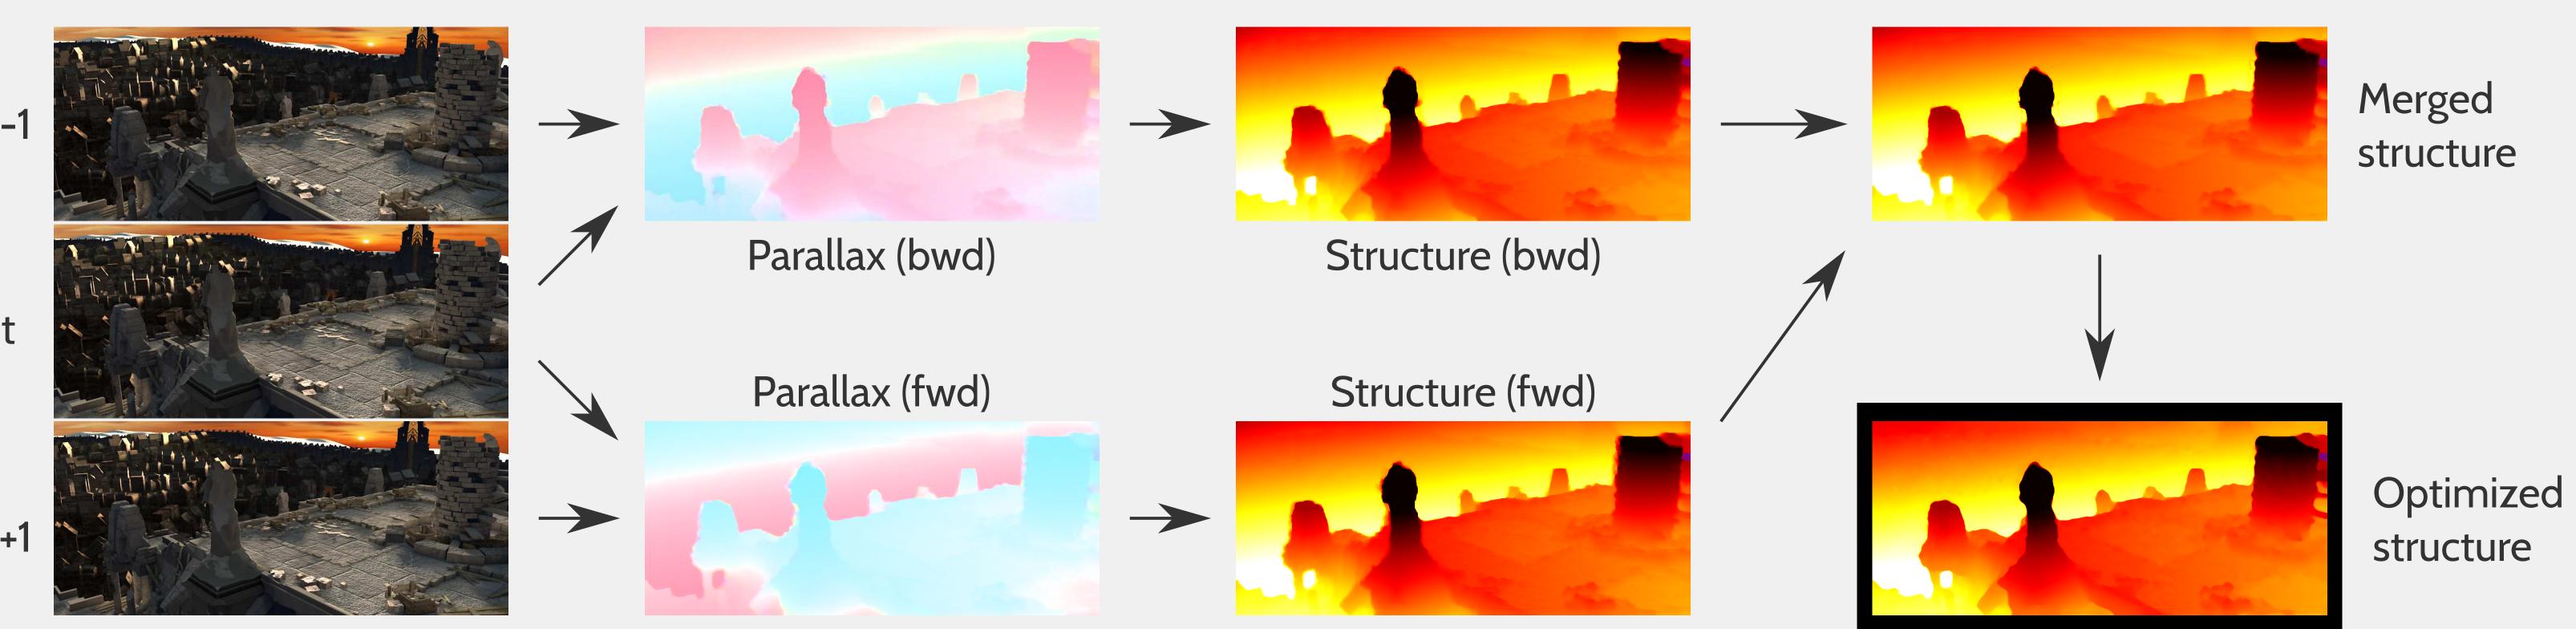
b(t): Camera motion. A single number per frame, varies in time.

A(x): Structure. A value per pixel, constant in time.

Depth of point x Distance 2nd camera to plane



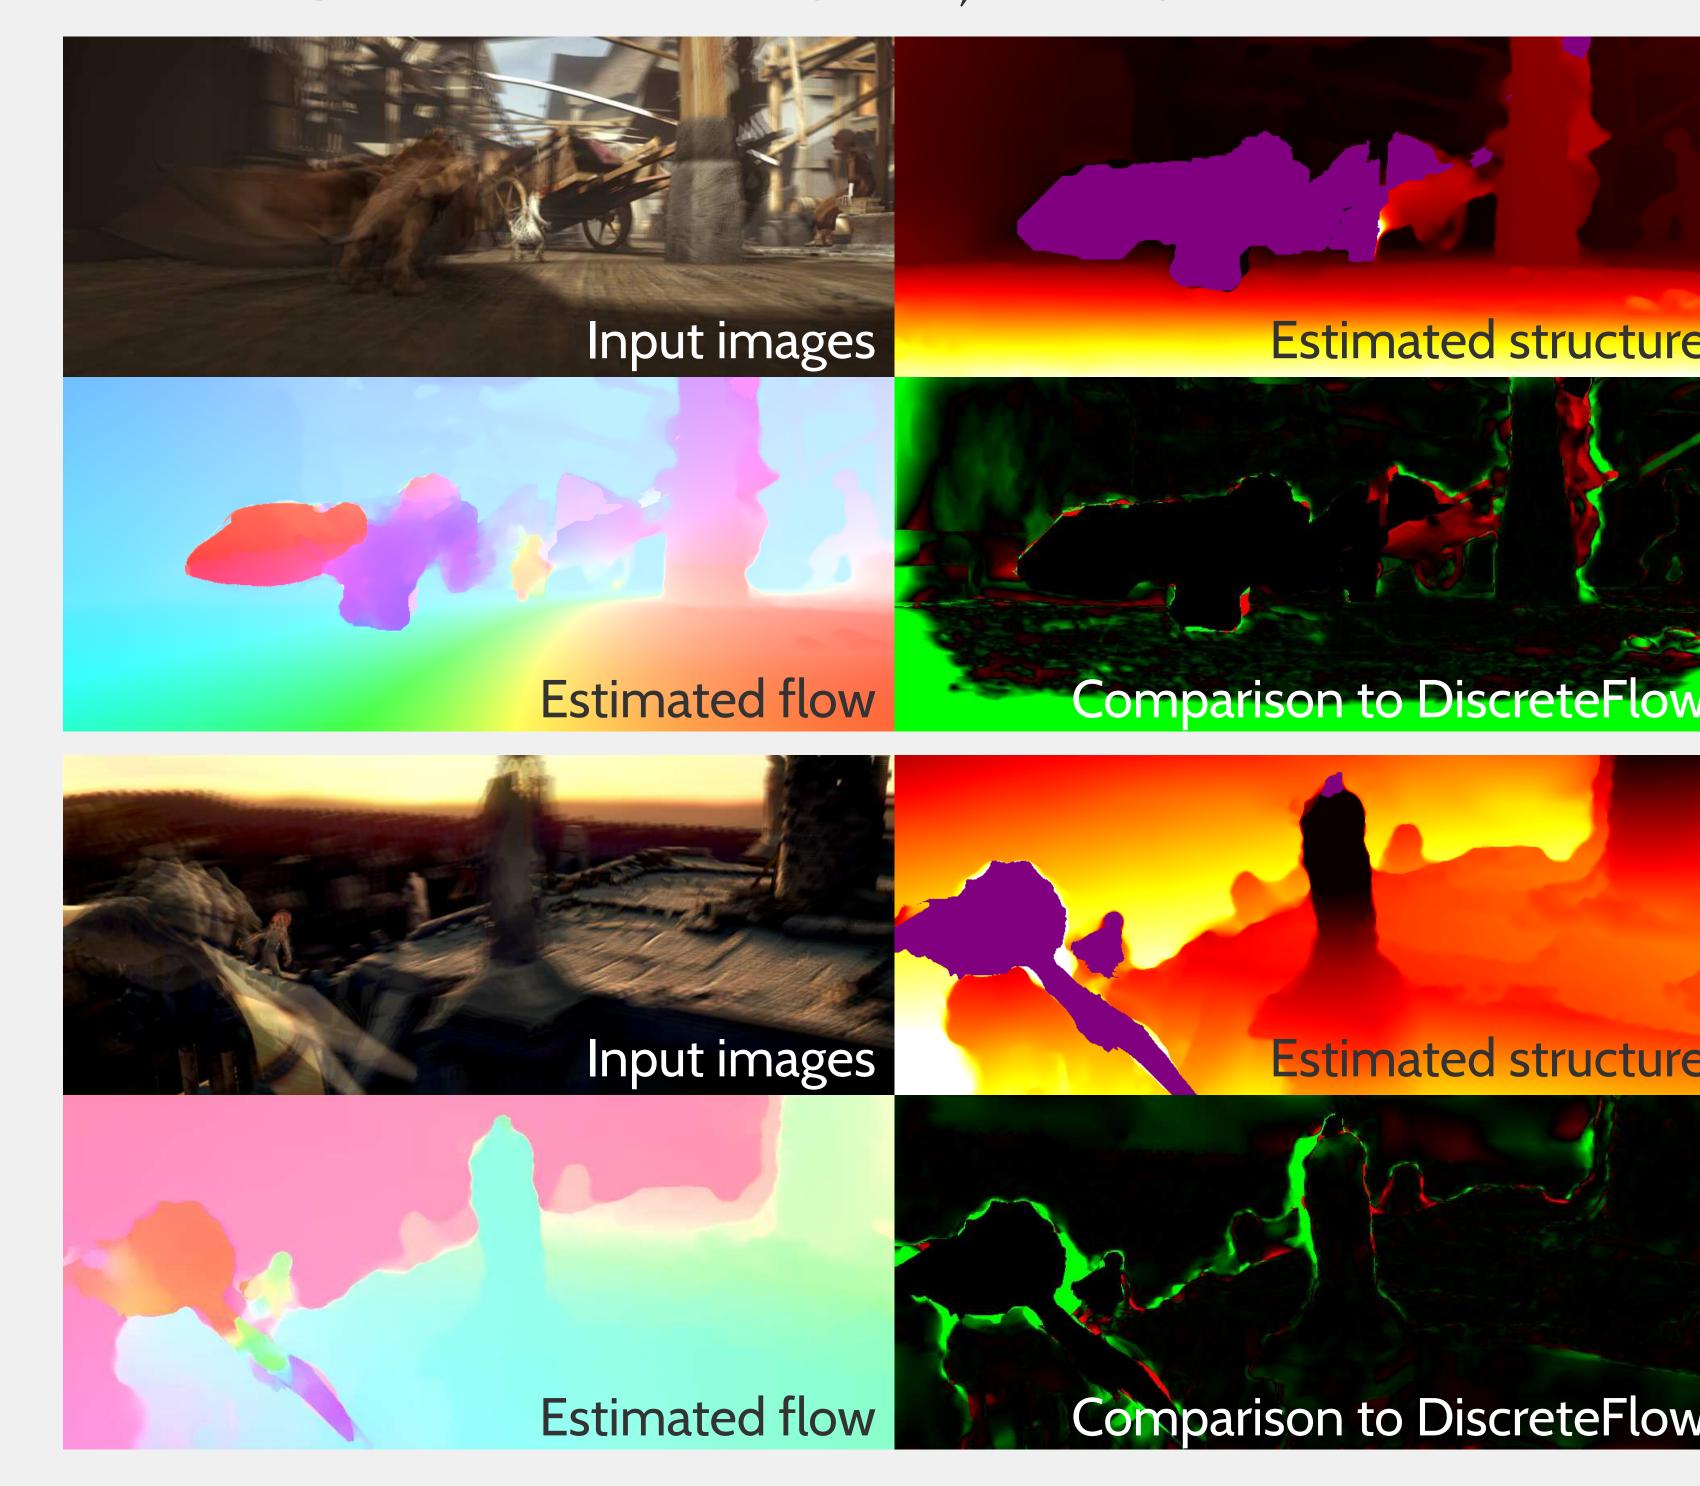
To improve results, especially in occlusions, we compute the structure in forward and backward directions and merge both, followed by a variational refinement.



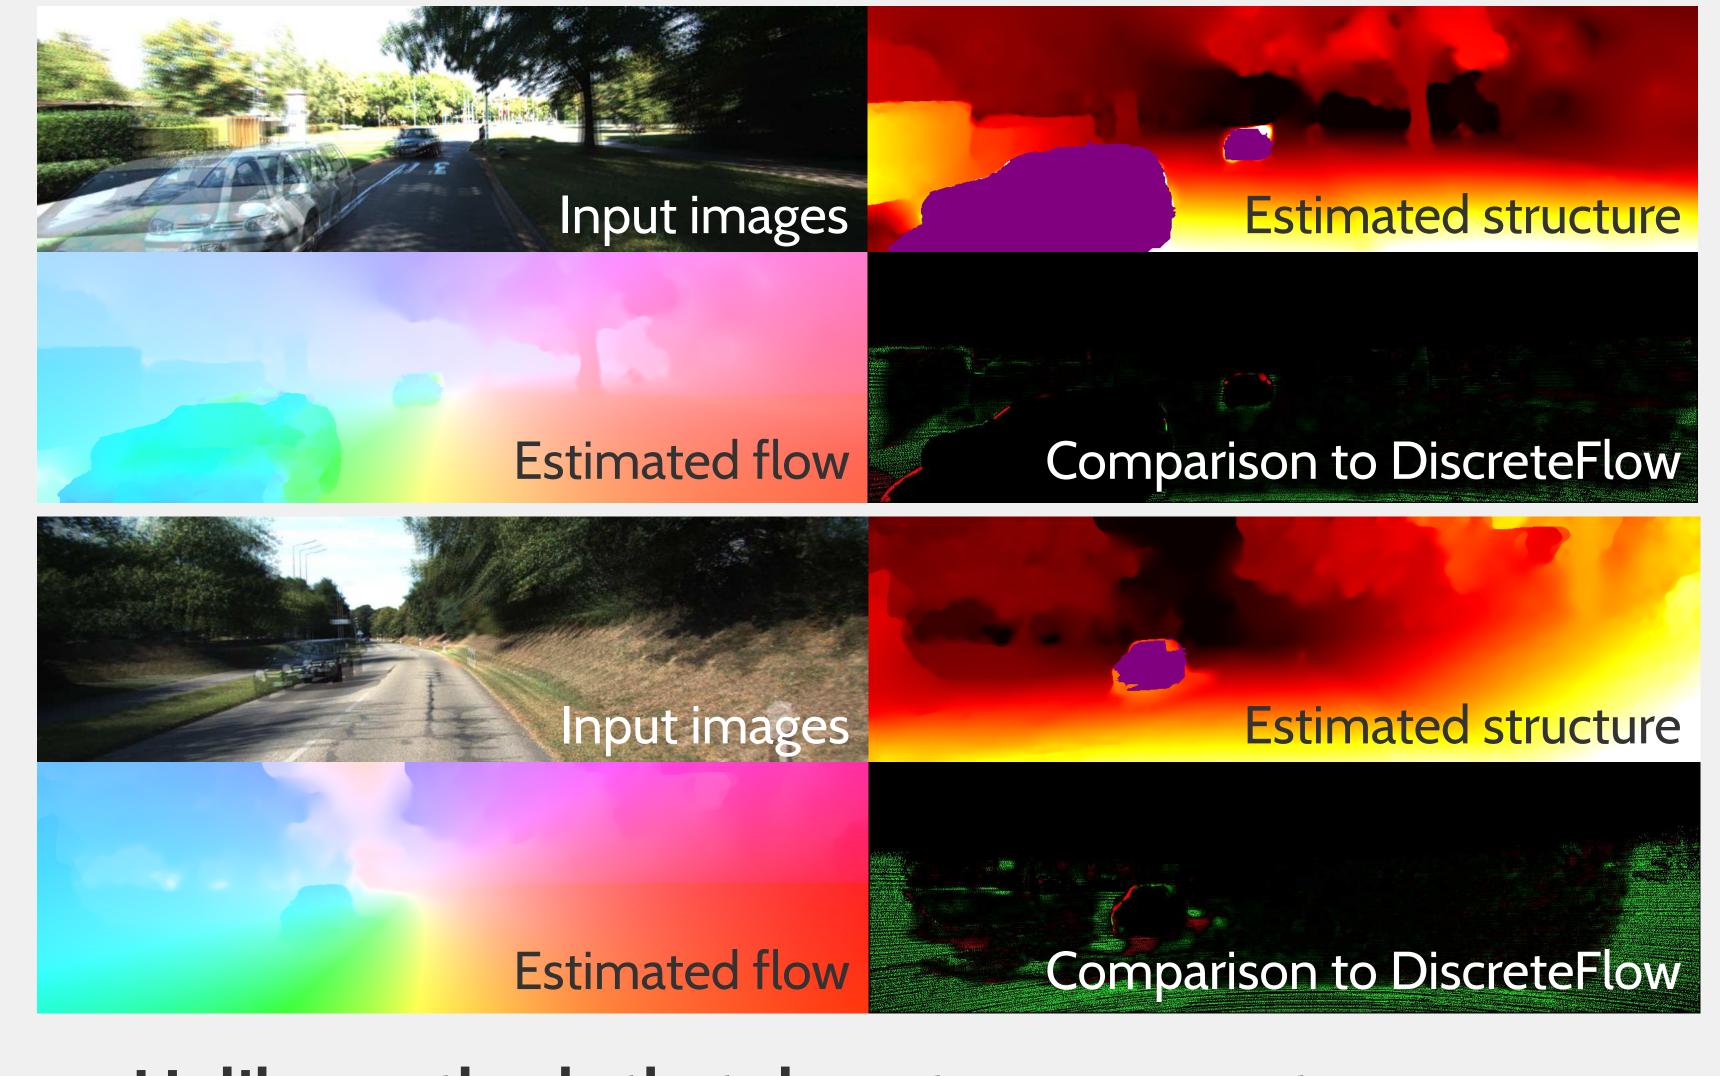
In moving regions, we use the results from DiscreteFlow.

Results

MPI-Sintel: Rank 1 on Clean, Rank 3 on Final.



KITTI: Rank 5 among monocular methods.



Unlike methods that do not use geometry, MR-Flow yields top results on both benchmarks.

Future work

- Combine geometry with end-to-end learning.
- > Stronger temporal consistency in segmentation.
- Higher-level reasoning (e.g. object support).