Large-Scale Semantic 3D Reconstruction: an Adaptive Multi-Resolution Model for Multi-Class Volumetric Labeling

M. Bláha, C. Vogel, A. Richard, J. D. Wegner, T. Pock, K. Schindler (CVPR 2016)

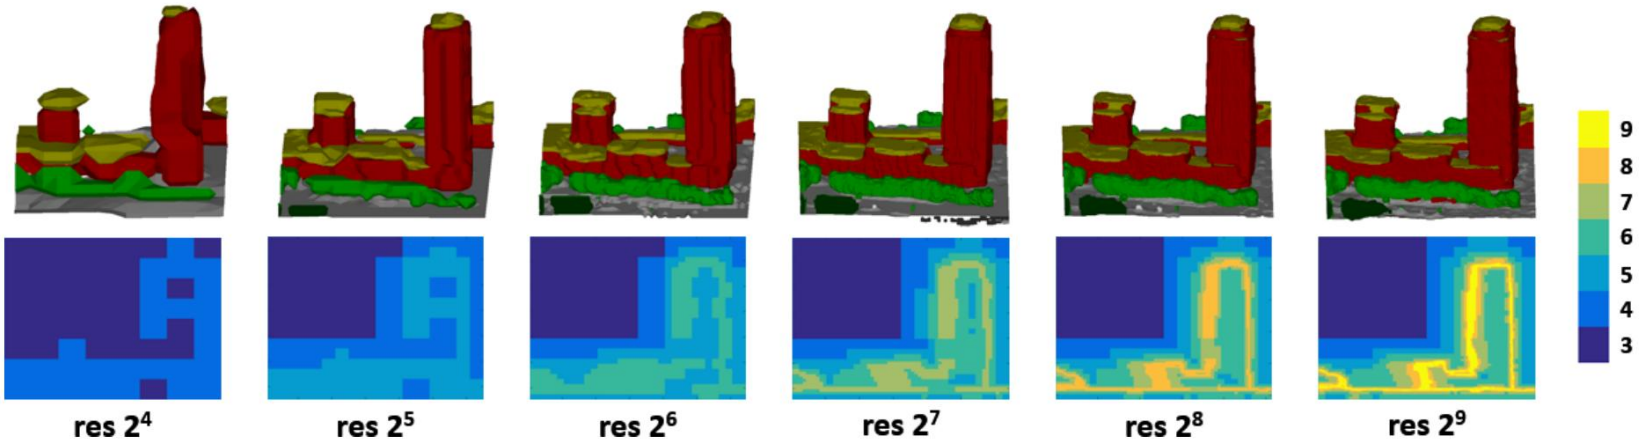

- Joint formulation of semantic segmentation and 3D reconstruction enables to use class-specific shape priors
- State-of-the-art could not scale to large scenes because of run time and memory
- Extension of an expensive volumetric approach
  - Hierarchical scheme using an Octree structure
  - Refines only in regions containing surfaces
  - Coarse-to-fine converges faster because of improved initial guesses
- Saves 95% computation time and 98% memory usage
- Evaluation on real world data set from the city of Enschede