TorontoCity: Seeing the World with a Million Eyes

S. Wang, M. Bai, G. Mattyus, H. Chu, W. Luo, B. Yang, J. Liang, J. Cheverie, S. Fidler, and R. Urtasun (ARXIV 2016)

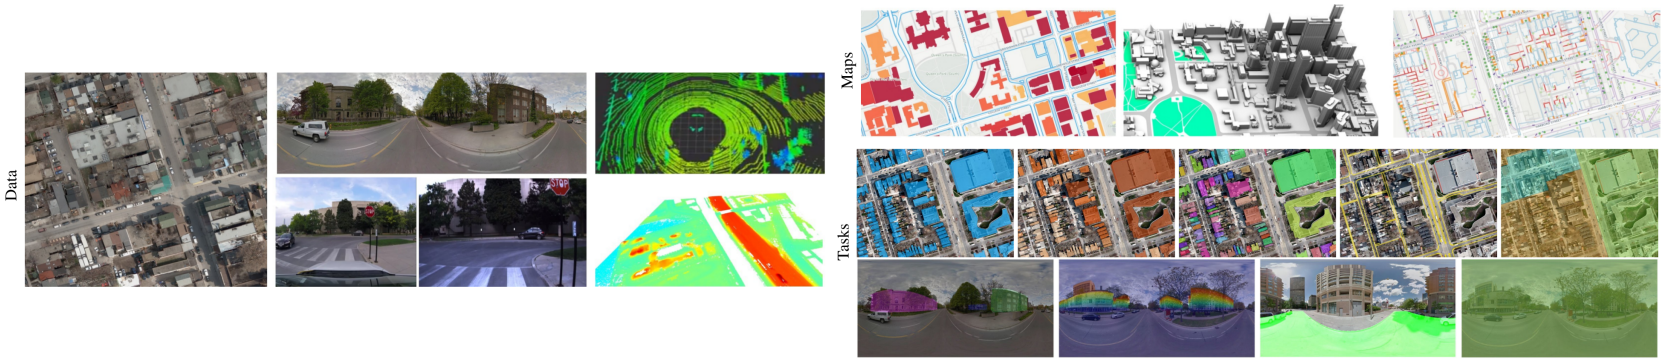

- Large-scale benchmark for multiple tasks covering full greater Toronto area (GTA) with 712.5km 2 of land, 8439km of road and around 400,000 buildings
- Limitations of current benchmarks:
  - Small set of sensors
  - Lack of rich semantics and 3D information at a large-scale
- Joint reasoning about geometry, grouping and semantics (three R's of computer vision)
- Captured from airplanes, drones and cars driving around the city
- Tasks: building height estimation (reconstruction), road centerline and curb extraction, building instance segmentation, building contour extraction (reorganization), semantic labeling and scene type classification (recognition)
- Utilizing different sources of high-precision maps to create ground truth
- Aligning all data sources with the maps while requiring minimal human supervision
- State-of-the-art methods work well on semantic segmentation and scene classification.
- Open challenges: instance segmentation, contour extraction and height estimation
- Many possible extensions