Learning Where To Classify In Multi-View Semantic Segmentation

H. Riemenschneider, A. Bodis-Szomoru, J. Weissenberg, L. V. Gool (ECCV 2014)

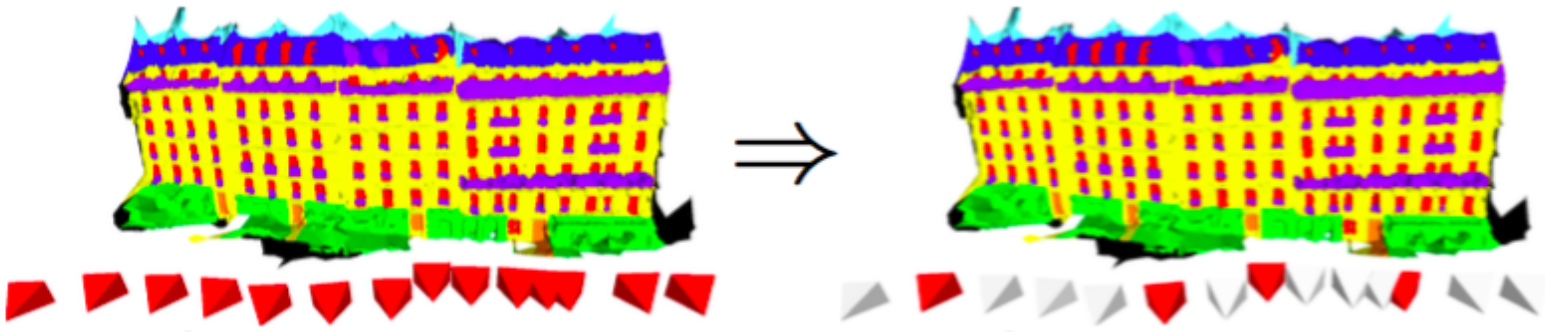

- View overlap is ignored by existing work in semantic scene labelling, and features in all views for all surface parts are extracted redundantly and expensively

- Contributions:
  - Proposes an alternative approach for multi-view semantic labelling, efficiently combining the geometry of the 3D model and the appearance of a single, appropriately chosen view - denoted as reducing view redundancy
  - Show the beneficial effect of reducing the initial labelling to a well-chosen subset of discriminative surface parts, and then using these labels to infer the labels of the remaining surface. This is denoted as scene coverage
  - Accelerates the labelling by two orders of magnitude and make a finer-grained labelling of large models (e.g. of cities) practically feasible
  - Provides a new 3D dataset of densely labelled images