**Using Semantic Relationship Among Objects for Geospatial Land Use Classification**

The current state of the art solutions to geospatial land classification utilize deep learning and cast the problem as many individual labeling problems. For each bounding boxed region, a label is sought that reflects one of the possible categories, including, sometimes a background category. In this work, we show that prior knowledge, in terms of semantic relationships among detected objects, allows us to leverage the semantics of the image in addition to visual features to obtain improvements in the recognition of objects. For instance, the presence of label road increases the chances of the label vehicle and vice versa.

The starting point consists of the top-k possible labels for each region, as detected by a deep learner. We use an ensemble, which we call Hydra, of twelve deep learning networks based on two state-of-the-art CNN architectures, ResNet and DenseNet. We use Grenander's canonical pattern theory formalism coupled with the commonsense knowledge base, ConceptNet, to impose context constraints on the labels detected by deep learning algorithms. The constraints are captured in a multi-graph representation involving generators and bonds with a flexible topology, unlike an MRF or Bayesian network, which have fixed structures. The minimization of the energy of this graph representation results in a graphical representation of the semantics in the given image.

We show our results on the recent IARPA Functional Map of the World (fMoW) dataset consisting of 1,047,691 images spanning 63 categories. We have the biggest improvement in performance was with respect to distinguishing the false detection class from the other 62 classes. About 67% of the regions that were fixed were either misclassified as false detection or misclassified as one of the other 62 classes when it was a false detection. The other categories with significantly improved performance were: zoo, nuclear power plant, flooded road, port, shipyard and border checkpoint. Overall, without context, the performance was 77.04% and with context, the performance improved to 77.98%. This performance is third best of 69 reported performance on the IARPA fMoW challenge.