Context Selective Decision Forest with applications to Lung Segmentation in CT

Team name: GE Global Research, Niskayuna NY

Albert Montillo

This is a short description. For full details see the full paper which has been submitted to the Pulmonary Image Analysis workshop at MICCAI 2011.

Overall structure of the algorithm is voxel classification. It is a fully automated algorithm requiring no input from the user.

The steps of the algorithm are:

1) preprocessing consisting of

   1a) normalizing intensities to Hounsfield units which includes ensuring that the corners of the scan are valid values no less than -1024 (replacing invalid values with closest valid ones),

   1b) normalizing orientation limited to one degree of freedom: upside down or not. This is based on the known volume orientation tag (e.g. RAI or RPI).

   This normalization does not otherwise adjust pose (rotation, translation) nor scale nor skew. We do this for 110 training examples.

2) Apply the trained decision forest classier to the never-seen-before LOLA11 challenge data.

3) Preprocessing of the test data is the same as the training data (normalize intensities to Hounsfield units and normalize upside-down or not orientation.

4) label relaxation

Limitations:

Currently a limitation of the CSDF is suboptimal performance on small diameter DFOV reconstructed images. A potential solution is to re-use our selective context idea (section 3) to limit the context used by the decision forest in the training data to exclude regions outside the cylinder that just encloses the lungs in the training data, because these regions are not reconstructed in small diameter DFOV images.

The training data

The training data consists of 110 CT scans of the full thoracic region from subjects of varying ages, weights, levels of health or pathologies and healthy subjects, from different scanners, different protocols (contrast and non contrast).

Average runtime of the algorithm

The parallel implementation of our voxel classification based method segments both lungs simultaneously in just 43 seconds per volume (where a typical volume is 256x256x250 at 2x downsampling) using a standard Intel Xeon 2.4GHz computer (8 core) with 16GB RAM running Win7 x64. A very good, coarse labeling (at 8x downsampling) can be achieved in 1.7 seconds. Training on the 110 volumes, which need only be done once, requires about 8 hours