MPI-Br-GE (1)

Authors: Coleman Broaddus, Martin Weigert, Uwe Schmidt, Gene Myers

Email: broaddus@mpi-cbg.de

Platform: Mac, Linux

Prerequisites: Python with standard scipy packages, tensorflow

MPI-Br-GE (1): SUMMARY

Motivated by making the best of weak annotations commonly available in practice, we propose to exploit nuclei centroid annotations. In particular, we train a convolutional neural network for centroid predictions of cell nuclei with substantially varying size that we use as accurate seed points for conventional pixel grouping methods (e.g. watershed). We perform pixel grouping on a nuclei probability image obtained from a trained classifier, thereby assuming some very sparse (and easy to obtain) pixel-wise annotations.

MPI-Br-GE (1): PREPROCESSING

Raw fluorescence images are normalized, linearly rescaling to set the 2nd and 99.5th percentile to 0 and 1 respectively.

MPI-Br-GE (1): DETECTION

We train a convolutional neural network to predict nucleus centroid locations by regressing a (Gaussian) kernel placed on each annotated center location. In particular we employ a 3D U-net that takes the 3D intensity image as input and predicts a single-channel output image of the same size. We choose an architecture with three resolution levels, employing two convolution layers of filter size 3×5×5 before each resolution change. We start with 32 convolution filters at the highest resolution, whose number is doubled (halved) after each pooling (upsampling) step. We use Adam to train the network with a mean squared error loss. After training, we extract the location of the centroids from the network prediction by finding local peaks via the *skimage.feature.peak_local_max* function, with *min_distance* set to the standard deviation of Gaussian kernel and *threshold_abs* fixed at 0.1.

MPI-Br-GE (1): SEGMENTATION

Given the predicted centroid locations from the detection CNN, we use these as seed points to perform a standard watershed segmentation to obtain an instance segmentation of all cell nuclei. Instead of using the (denoised/blurred) image intensities directly as the potential for the watershed, we instead use a nuclei probability image obtained from a pixel-wise classifier.

Since we assume only very little pixel-level segmentation annotation, we train a very small CNN with relatively few parameters. Specifically, we only use two 3D convolution layers with 16 filters each and kernel size 3×7×7. The final convolution layer uses a single kernel of size 1×1×1 and is followed by a sigmoid activation. We again use Adam for training the network with a binary cross-entropy loss.

MPI-Br-GE (1): POST-PROCESSING

No post-processing is carried out after segmentation.