UC San Diego
Jacobs School of Engineering

Biomedical Image Segmentation

Lucas Tindall, Amir Persekian, Max Jiao

Abstract

Traditional CNN architectures have been very good at tasks such as applying a class label to an image after training on a very large example dataset. However, in many applications, a more fine grained classification is necessary. One example is for image segmentation in the Biomedical field. In order to classify pixels in images based on only the contextual information surrounding it, without a need for a large training dataset, we use an efficient and generalizable network called the U-net. The U-network leverages the homogenous nature of medical images to train on similar images very efficiently. Our network is able to classify with an F1 score of 0.96 after training on just 10% of our images.

Datasets

- Kaggle Dataset[2]
  - Identify nuclei in cells
- NCMIR Dataset[3]
  - Identify mitochondria
- ISBI Datasets[4][5]
  - (2012 Challenge) Segmentation of Neuron structures in EM stacks
  - (2013 Challenge) 3D segmentation of neurites in EM images
- EPFL CVLab[6]
  - Electron Microscopy Dataset

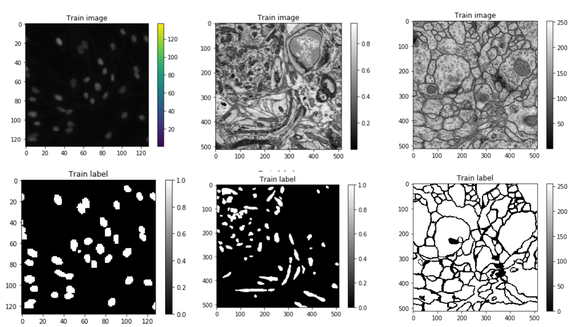

Methods

Preprocessing

- Data augmentation (random crops, rotations, affine shifts, horizontal/vertical flips)
- Small datasets with homogenous features benefit from more varied training samples

Model Architecture

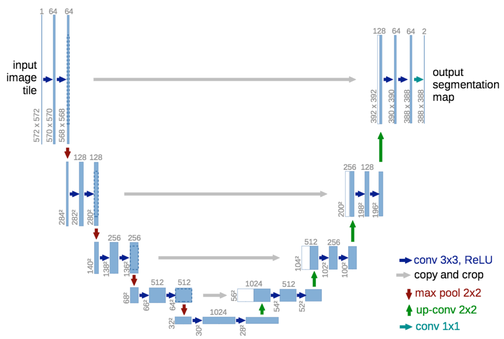

The Unet[1]

- First half is a traditional convolutional network with max pooling
- Second half is a deconvolutional network with skip connections
- Mini-batch Stochastic Gradient Descent using Adam optimization

$$BCE\ loss = -(y * log(p) + (1 - y) * log(1 - p))$$

- Binary Cross Entropy loss
- Batch size varied from 1 – 32
- Learning rate of 0.01

Postprocessing

- Instance segmentation requires different labels to distinguish each object in a mask
- K-means to find the centers of nuclei using elbow method to find optimal K value
- Use markers from K-means to perform marker-controlled watershed segmentation

Results

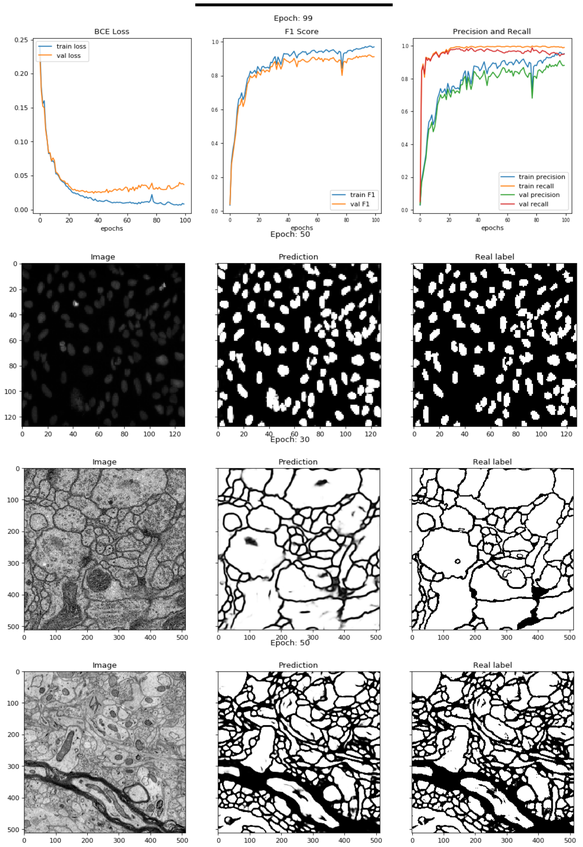

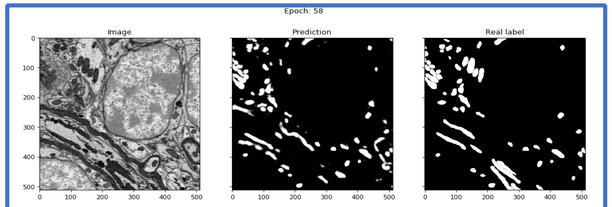

Marker-Controlled Watershed

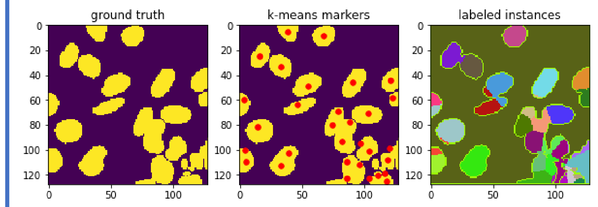

F1 Scores

Kaggle	NCMIR	ISBI 2012	ISBI 2013	EPFL
0.9	0.89	0.93	0.96	0.91

Conclusion

It can be concluded that the U-net designed by [1] is very effective at processing contextual data and generalizing it through the features of similar images. However, there are certain limitations in that the network cannot be easily generalizable across multiple different datasets at the same time; the system does not perform well to images that are too unfamiliar, and it is not easy to gauge the tolerances for data augmentation for different datasets. Postprocessing techniques can to an extent distinguish components on generated masks, but the network itself does not have the capability to. These are areas that can be further explored in future work.

References

[1] Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. "U-net: Convolutional networks for biomedical image segmentation." In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234-241. Springer, Cham, 2015.

[2] https://www.kaggle.com/c/data-science-bowl-2018#description

[3] https://ncmir.ucsd.edu

[4] http://brainiac2.mit.edu/isbi_challenge/home

[5] http://brainiac2.mit.edu/SNEMI3D/home

[6] https://cvlab.epfl.ch/data/em