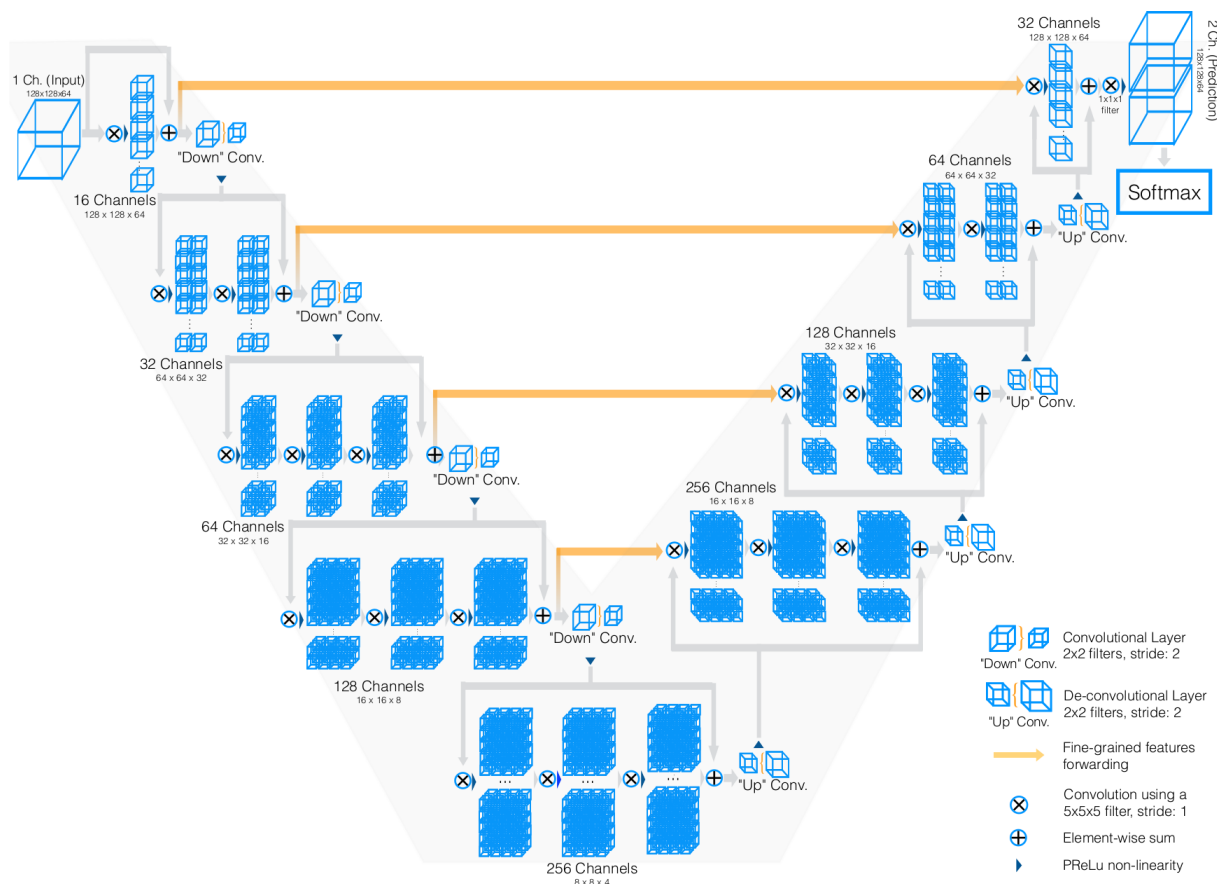


## A PyTorch implementation of V-Net



Vnet is a PyTorch implementation of the paper V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation by Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. Although this implementation is still a work in progress, I'm seeing a respectable 0.355% test error rate and a Dice coefficient of .9825 segmenting lungs from the LUNA16 data set after 249 epochs. The official implementation is available in the [faustomilletari/VNet](#) repo on GitHub.



This implementation relies on the LUNA16 loader and dice loss function from the Torchbiomed package.

### Differences with the official version

This version uses batch normalization and dropout. Lung volumes in CTs are ~10% of the scan volume - a not too unreasonable class balance. For this particular test application I've added the option of using NLLoss instead of the Dice Coefficient.



### What does the PyTorch compute graph of Vnet look like?

You can see the compute graph here, which I created with `make_graph.py`, which I copied from `densenet.pytorch` which in turn was copied from Adam Paszke's gist.

### Credits

The `train.py` script was derived from the one in the `densenet.pytorch` repo.