10.1148/ryai.2019190015

Appendix E1

Section 1: Learning Rate Selection

Investigating the impact of image resolution choice on machine learning performance in radiology applications requires understanding the consequences of other relevant hyperparameter selections and modeling choices. One important hyperparameter is the selection of the learning rate which corresponds to the step size of the optimization process related to the calculation of the loss function's gradient. Figure E1 of our supplement shows learning rate finder curves for the chest x-ray label of "emphysema" for multiple choices of image resolution. Ideally, we would like to select the largest learning rate possible which also gives us low validation set loss, and we would preferentially underestimate the optimal learning rate rather than overestimate it to achieve best possible performance. Based on examining these learning rate curves for multiple diagnosis labels, we decided to fix our learning rate at 0.0005 for all the binary classification networks that we trained. For comparison, this learning rate is within an order of magnitude of the 0.001 learning rate in (2) and the 0.0001 learning rate in (3) for chest x-ray deep learning. Notably, though, 0.0005 is a particularly aggressive choice based on our learning rate search, and 0.0001 or 0.00001 would likely be more conservative or appropriate choices for achieving maximum possible cross-validation performance.

Section 2: Receiver Operating Characteristic Examples

While examining area under the receiver operating characteristic is often a useful way to compare performance between multiple models, ROC plots themselves are useful for the more granular information that they provide. Figure E2 provides 9 example ROC curves for 3 diagnosis labels (edema, mass, and nodule) for ResNet34 models trained for 3 subsample epochs. Input image resolution for these models is 256 × 256, 320 × 320, and 448 × 448.