Deep Learning for Chest X-Rays (CXRs)

Christine Tataru, Darvin Yi, Archana Senoyas, Anthony Ma

Patient with suspected lung disease → Primary care doctor orders chest x-ray → ★ Diagnosis by radiologist → Further tests for high risk patients

The starred step can take days or weeks! My goal is to create an automated triage system that detects abnormality in chest x-rays to help radiologist do their job more efficiently and primary care physicians proceed down this pipeline faster.

Data

- Stanford Radiology
- 400,000 CXRs (50,000 used for this project)
- Labeled normal, abnormal, and very high risk by expert radiologist

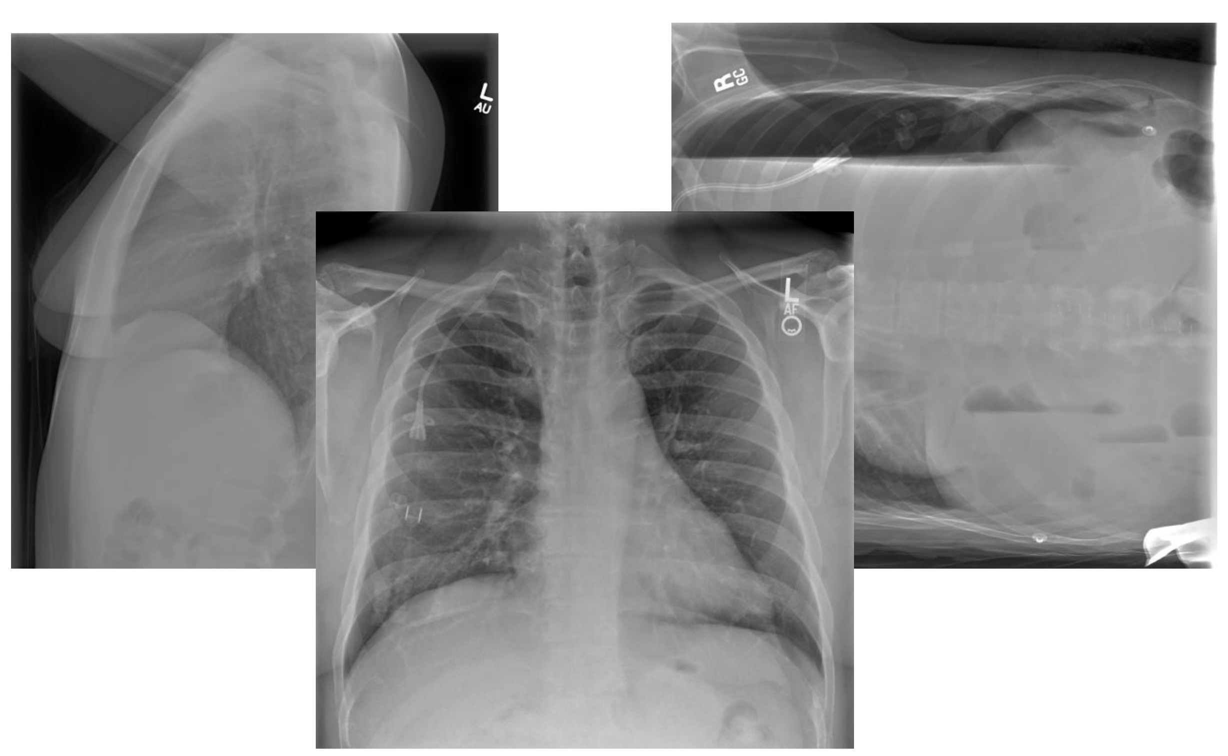

- A lot of variance in contrast, rotation/position, and image views

Methods

- **Histogram equalization** to increase contrast

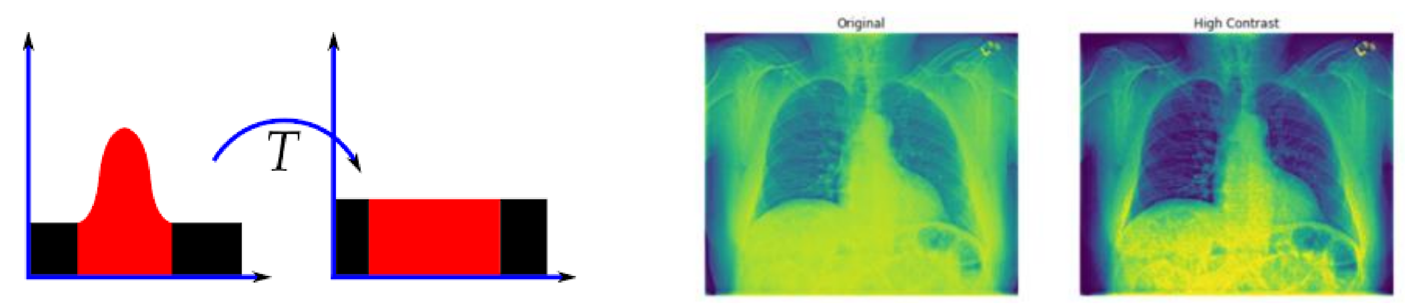

- **Down-sampling of images** from 3000 x 3000 pixels to 512x512 and 224x224
- **Data augmentation to prevent overfitting**
  - Randomly rotated 90, 180, 270, or 0 deg.
  - Randomly flipped left to right (or not)
  - Augmented with gaussian noise
- **Neural Net Architectures**
  - GoogLeNet: complex but not overly burdened with parameters
  - InceptionNet: more complex, same intuition
  - ResNet: very deep with different connectivity scheme

Results

- **Complexity of the architecture does not improve model performance** on 224px images (around 0.4 loss, random is 1.09)

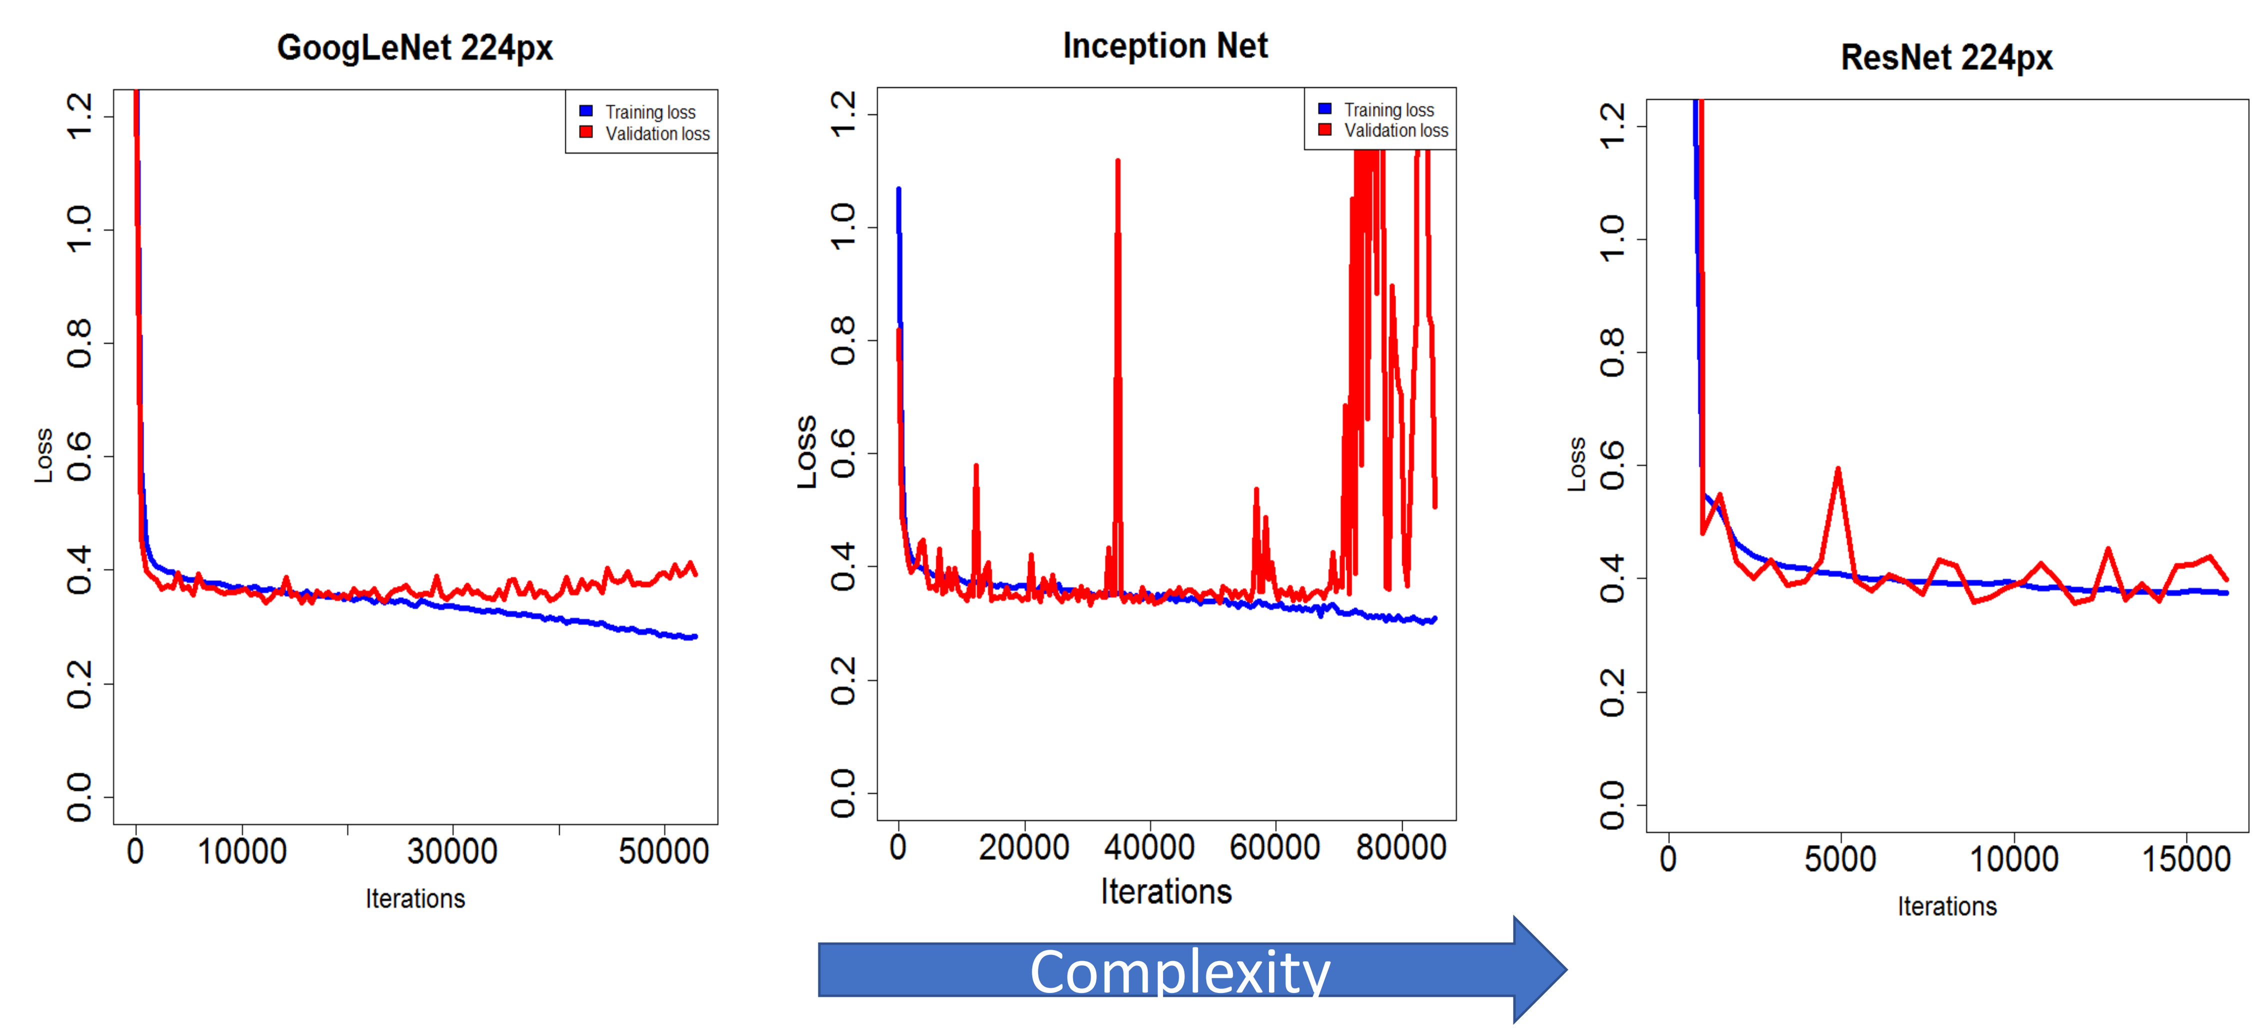

- **Using larger (512px) images does not improve performance** with GoogLeNet (again 0.4 loss, random is 1.09)
- **GoogleNet performs better with more data, up to 50,000 images**

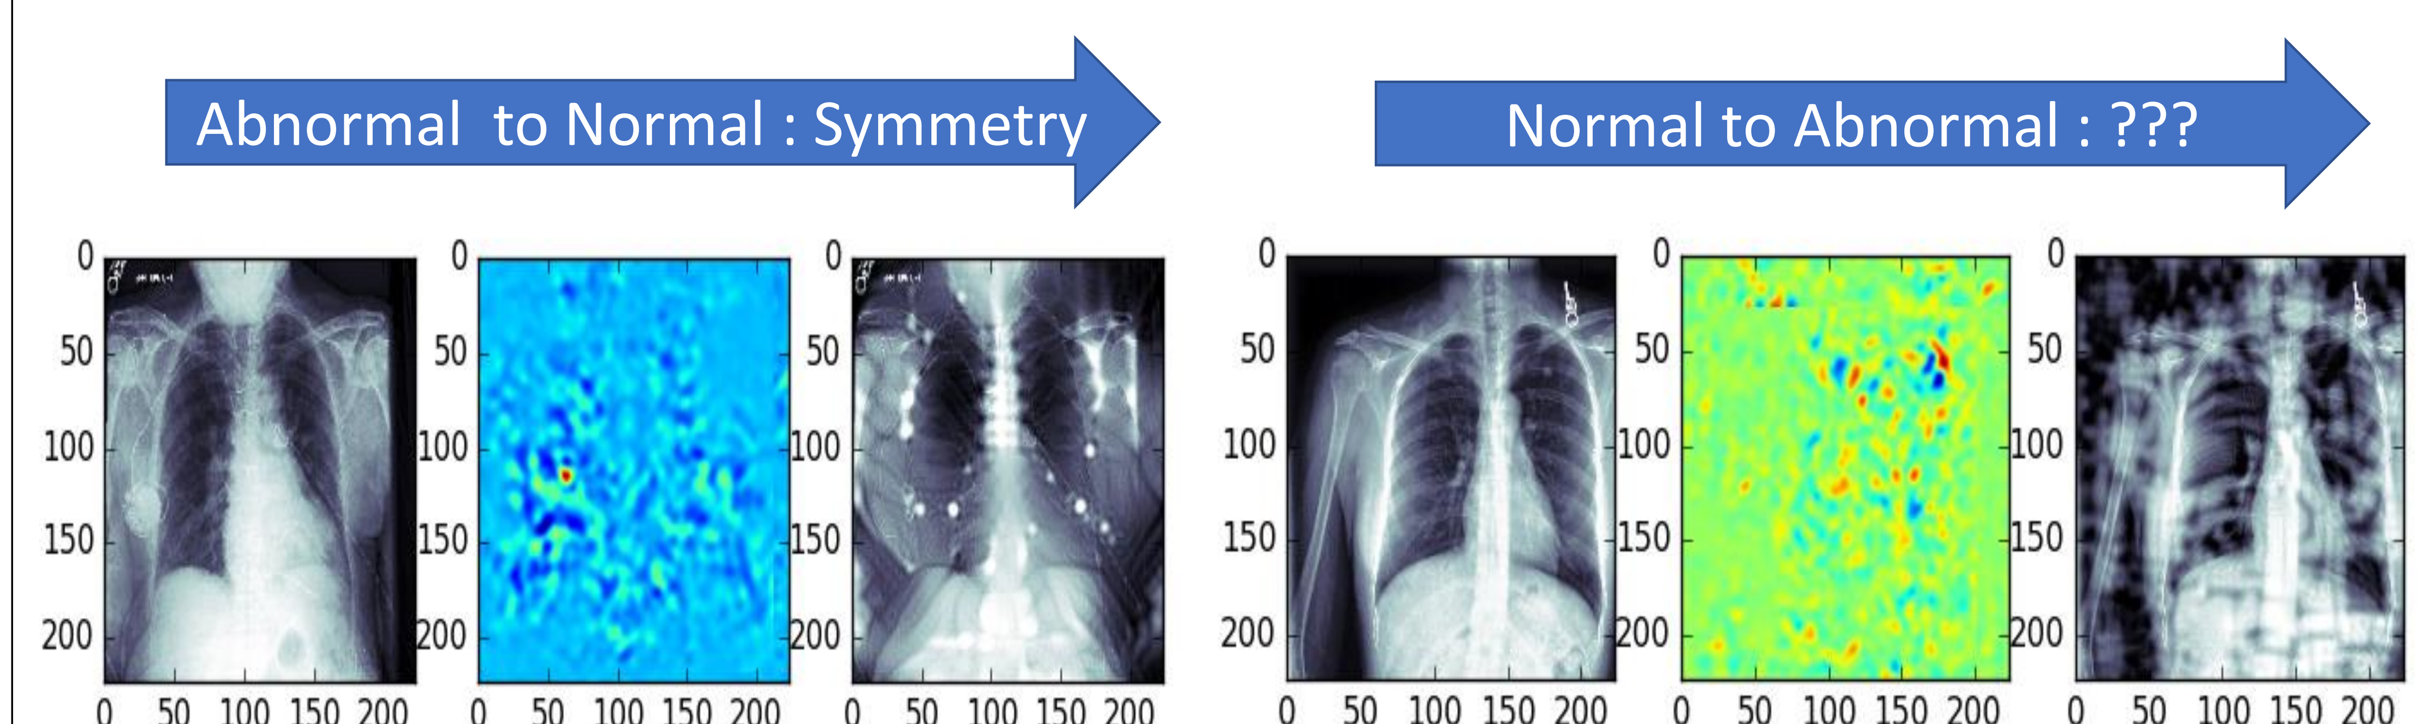

Conclusion

- The best performing CNN model has loss ~0.35 for 3-ary classification
- The model is robust to image size and architecture
- Visualization shows that macroscopic features are learned effectively by the model
- Future directions include further preprocessing (cropping lungs from images, cropping edges) and adding attention layer to the model.