Innovation/Impact: To the best of our knowledge, we are the first to use GAN-based anomaly detectors for active learning. We construct an active learning paradigm that is more accurate than current paradigms in the field. Our methods can be used to diversify medical imaging datasets, prioritize images for labeling, and detect images on which a trained segmentation model will likely fail.

Motivation: In Figure 1, we see examples of where a trained liver segmentation model failed when given new input that contained salient features. The liver segmentation model had a Dice coefficient of 97.1 on test data. The purpose of our model is to detect such “anomalous” images.

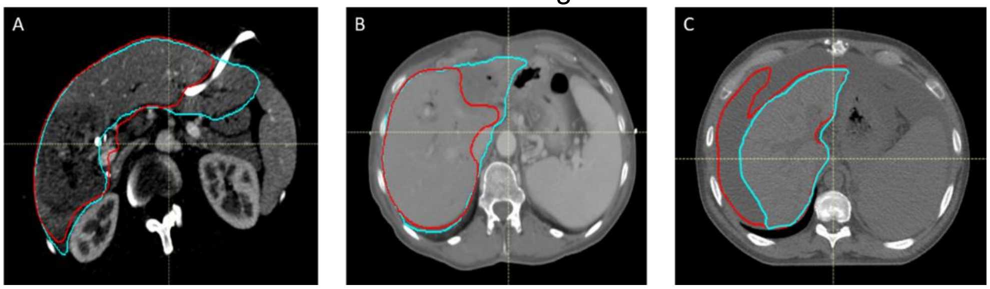

Figure 1: Instances where a trained liver segmentation model failed due to abnormalities that the model did not see during training. (A) stent (B) tumor (C) ascites.

Key Results: In Figure 2, we present images generated using our generator from a trained StyleGAN2 network. This is the generator that we use in our anomaly detection scheme. Figure 3 shows reconstructed images. Our generated images achieved a Fréchet Inception Distance of 31.102. In the pixel distribution space, it achieved a Wasserstein distance of 1.388, KL-divergence of 0.004, and mutual information of 0.241.

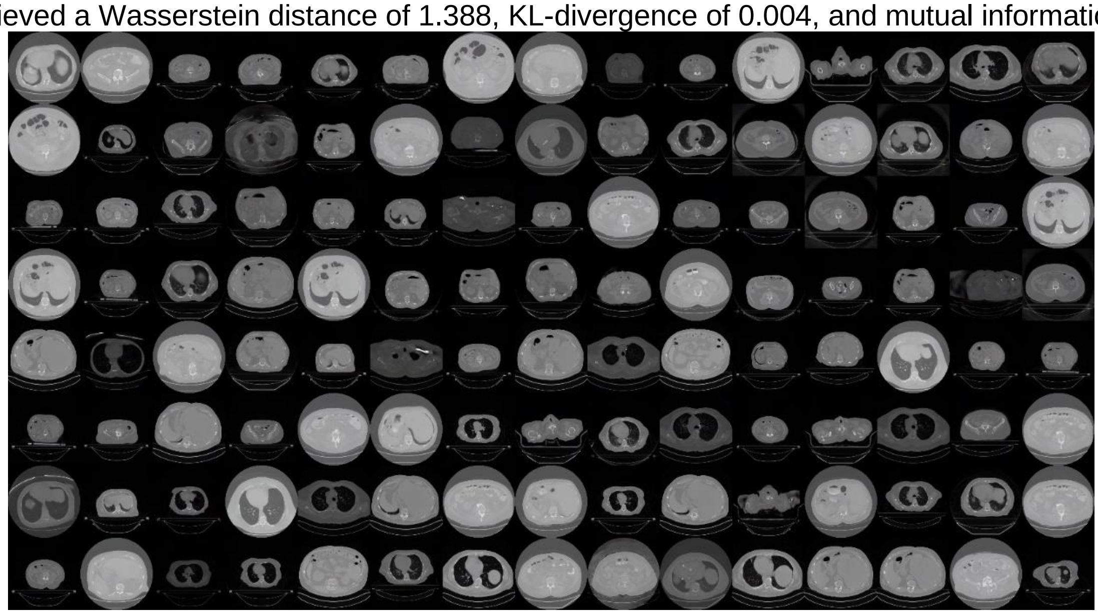

Figure 2: Generated slices of CT scans. Individual images are 512x512 but were reduced in size to be shown together. Images were randomly selected.

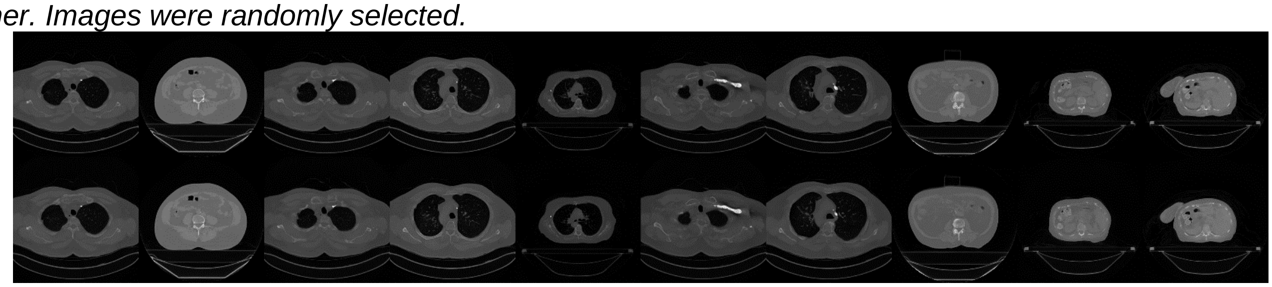

Figure 4: Top-10 performing reconstructed images. The top row contains the original images. The bottom row contains the corresponding reconstructed images. We randomly selected either a fake or a real image from each pair and gave them to a radiologist who scored both a specificity and sensitivity of .6.

Discussion: The presented results are the first steps in creating a robust anomaly detection scheme. From the reconstructed images, we can create an anomaly detection score. Preliminary scores yielded significant noise. Further work will reduce this noise via conditioning the GAN on relative slice position and training an encoder network to map images to the latent space. Once the noise in reconstructing images within the training distribution is resolved, our anomaly score will reflect when abnormalities are present in the novel data.