SAND2019-9331C

Deep learning for automated defect detection in high-reliability electronic parts

PRESENTED BY

Emily Donahue

Authors: Emily Donahue, Kevin Potter, Cari Martinez, Matthew Smith, Tu-Thach Quach, Christian Turner

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

X-Ray CT scanning is indispensable for component analysis

At Sandia and partner sites, thousands of components are imaged with X-Ray computed tomography (CT)

Non-destructive imaging allows detection of potential defects

Current QA largely depends on human inspection, as commercial software has not been developed for anomaly detection in this space

- 1000's of CT slices are manually inspected for conformity with product specification
- Can take days to analyze only a few centimeter-scale parts
- Costs millions of dollars and months of effort to diagnose component issues in testing

The 3D, non-destructive, high-resolution imaging that CT offers comes with a high analysis price tag

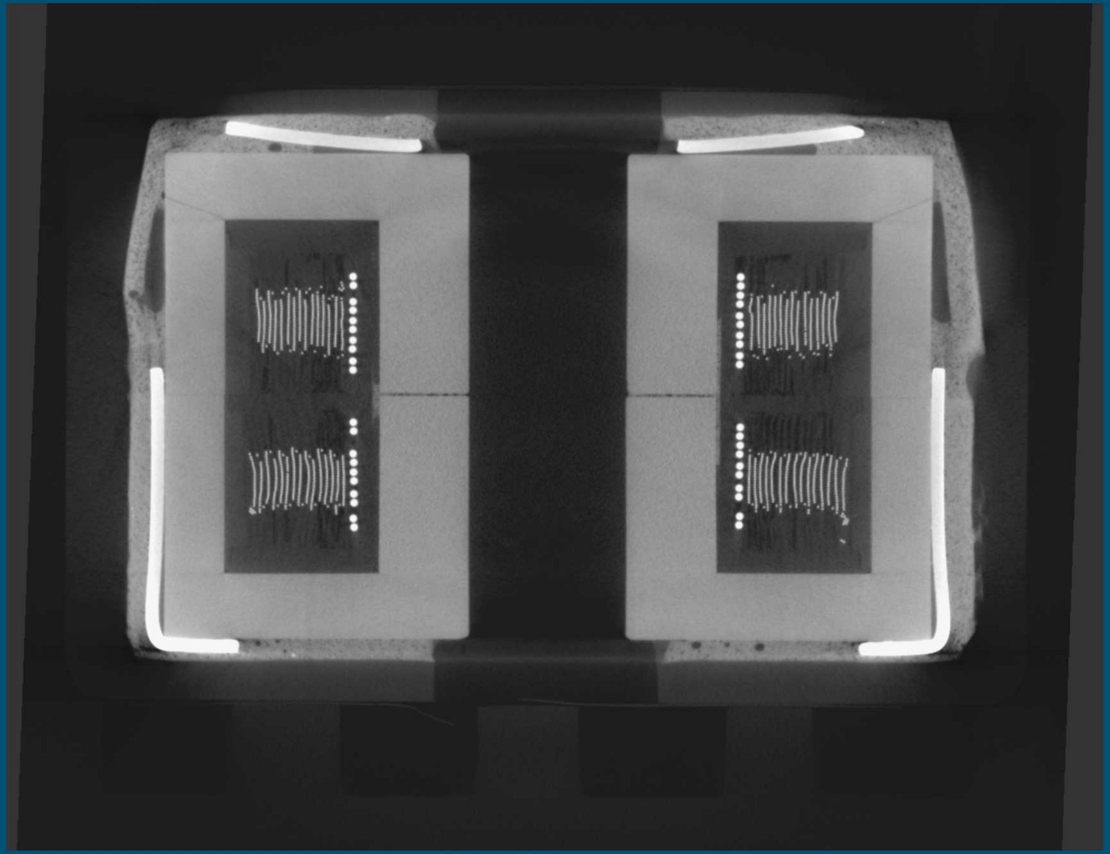

Cross-section of a magnetic flyback from X-Ray CT scan

Recent advances in generative techniques hold the answer

Generative models produce more and more realistic results every year

Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) have recently been applied to anomaly detection in various domains

- Retinal scans
- Cyber-physical systems
- Web traffic analysis

Faces generated by StyleGAN. Karras et. al. "A Style-Based Generator Architecture for Generative Adversarial Networks"

Category-conditioned images from BigGAN Brock et. al. "Large Scale GAN Training for High Fidelity Natural Image Synthesis

Training Generative Adversarial Networks (GANs)

System of 2 neural networks

- Generator: objective function is to create images that the discriminator will label as "real"
- Discriminator: objective function is to determine if an image is real or was generated

Adversarial training scheme produces a generator that has learned to model a distribution over natural images

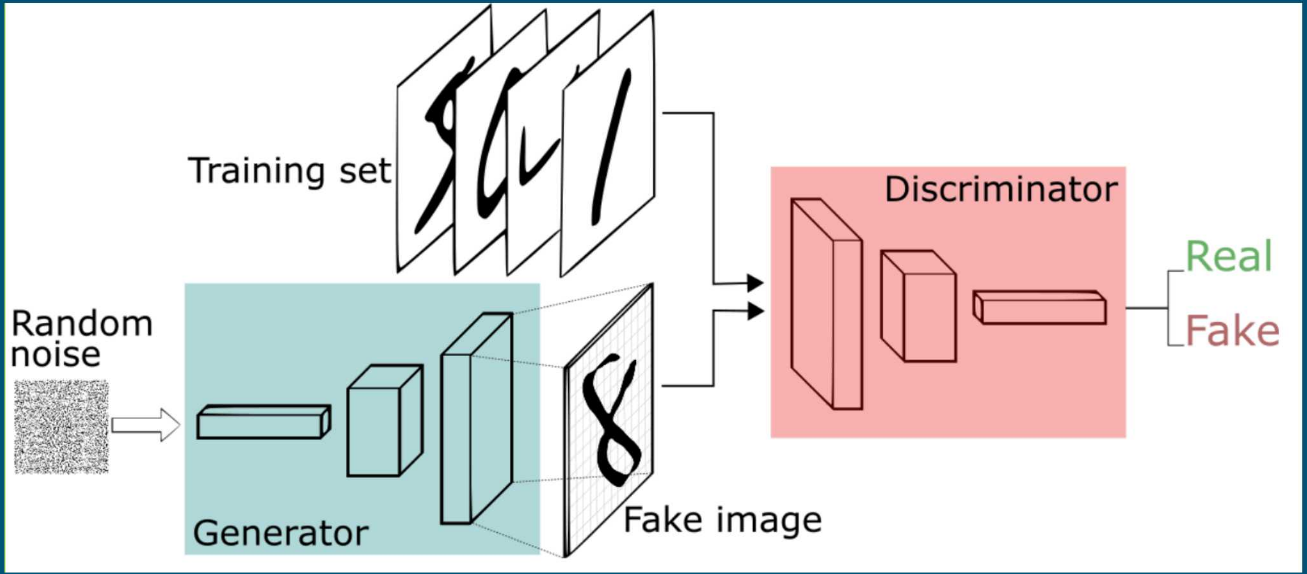

Schematic of GAN training
From skymind.ai:A Beginner's Guide to Generative Adversarial Networks (GANs)

The AnoGAN is an unsupervised deep learning system that can be trained for anomaly detection

The generator is trained on only normal (i.e. non-anomalous) images and learns to model a distribution over them

Inference involves generating a similar-looking image to a query image, and taking the difference to highlight anomalies

- Begin with a random latent vector, z_0
- Forward pass through generator to get an image, $G(z_0)$
- Calculate loss w.r.t query image, $L(x_q, z_0)$
- Backprop to get new latent vector, z_1
- Repeat for $\Gamma = 500$ steps

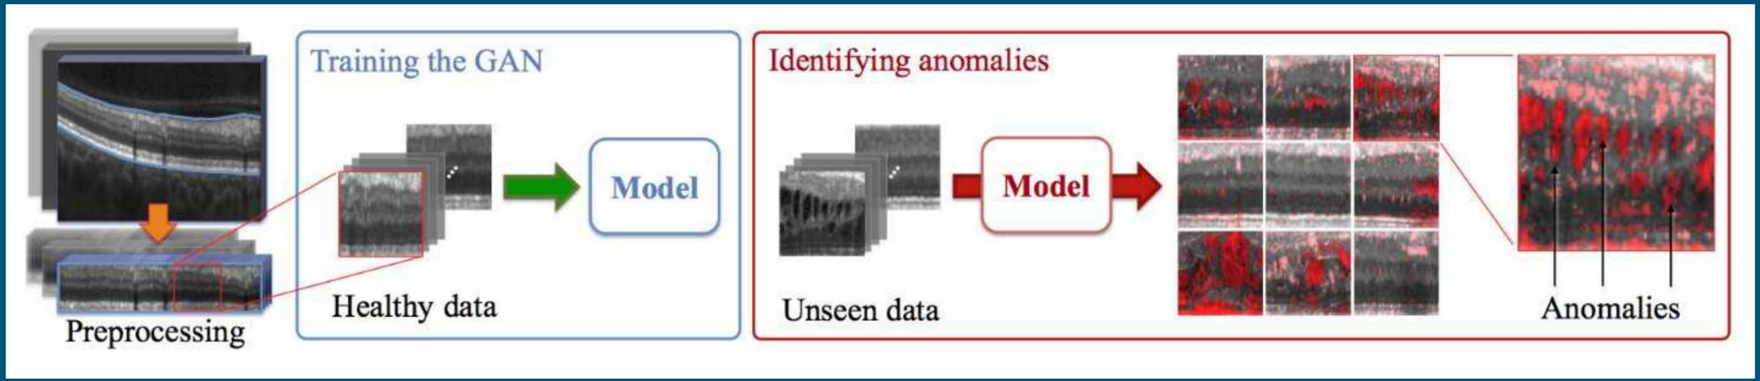

Overview of AnoGAN system from Schlegl et. al. "Unsupervised Anomaly Detection with Generative Adversarial Networks to Guide Marker Discovery"

We applied the AnoGAN method to CT scans of electronic components

Extracted regions known to have issues from scans of current viewing transformers and capacitors

Used axis-aligned CT scans of current sense transformers and capacitors

- Extracted regions-of-interest using known geometry of components
- Used traditional, statistics-based computer vision techniques

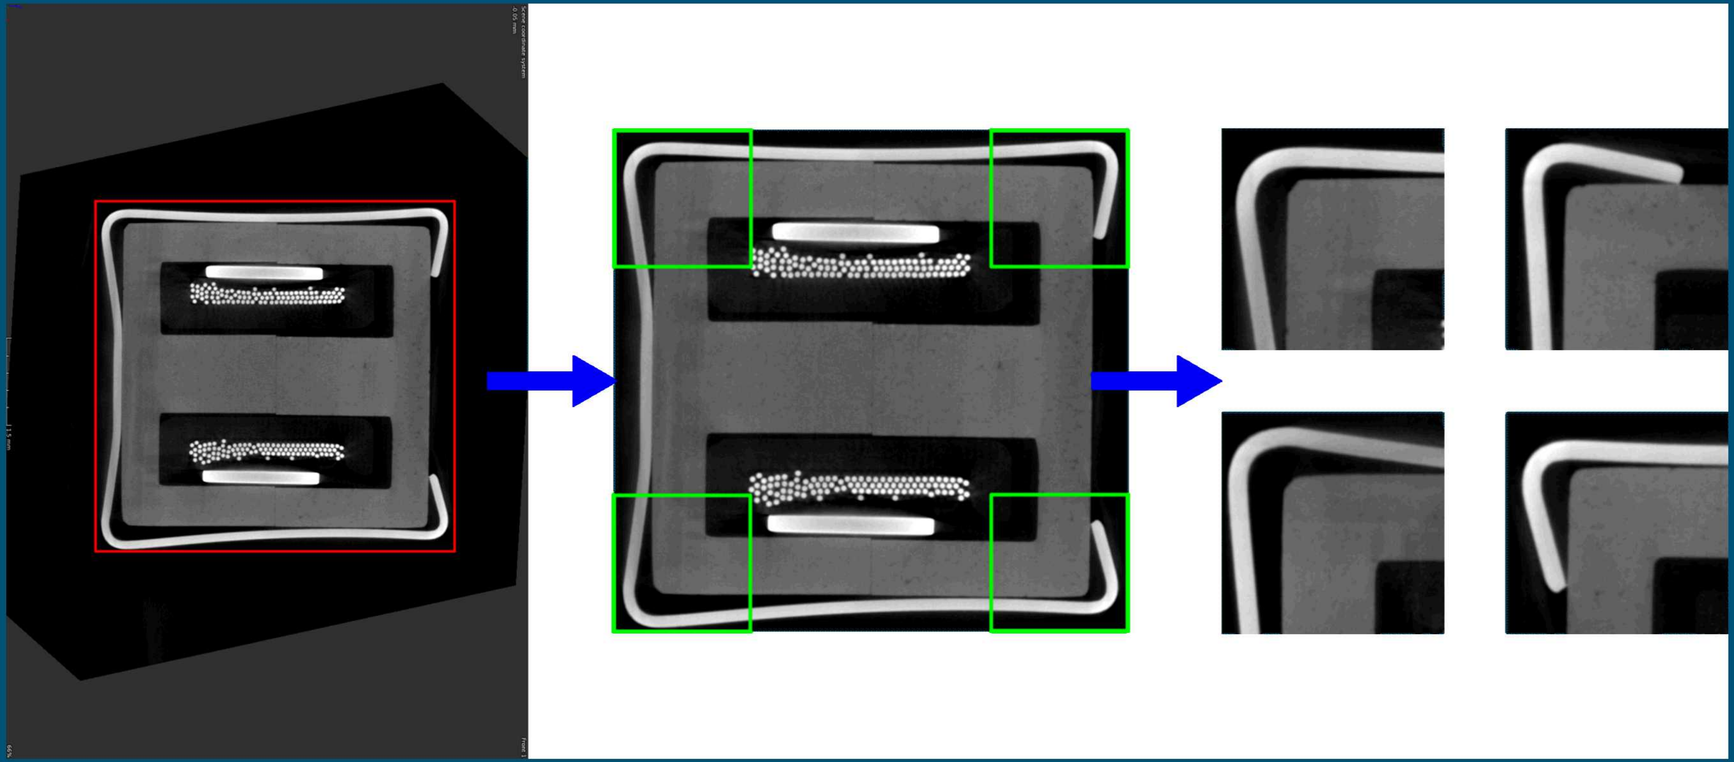

Qualitative results are promising for current transformer

Real query image

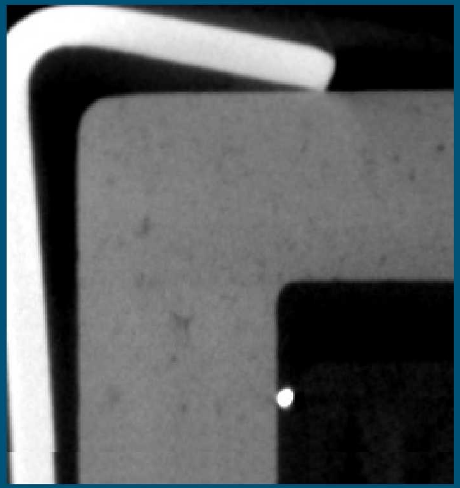

Generated image

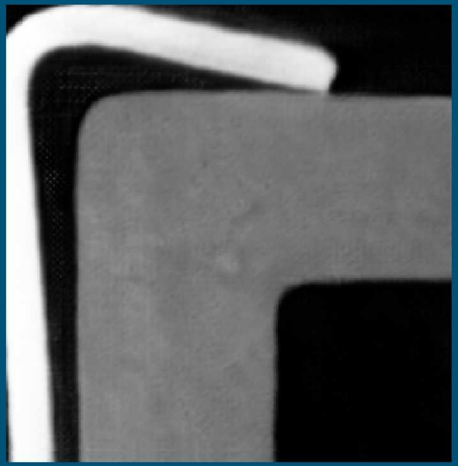

Anomaly detection

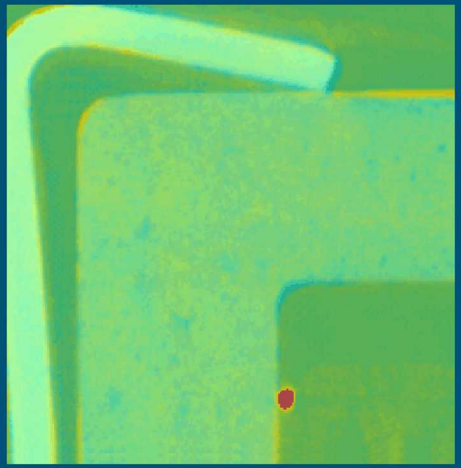

Undesirable solder ball highlighted in red

Training details

- Trained on 2996 non-defective part images
- 70 epochs
- 3.25 hours to train, ~10 seconds for inference
- 2 convolutional layered discriminator
- 2 de-convolutional layered generator

How does the latent dimension affect reconstruction?

Latent dimension (Λ): dimensionality of the random vector input to the generator

- The generator has to learn a mapping function $f\colon \mathbb{R}^{\Lambda} \to \mathbb{R}^{256\times256}$
- A larger value of Λ allows a more complex mapping to be learned, more degrees of freedom for the domain

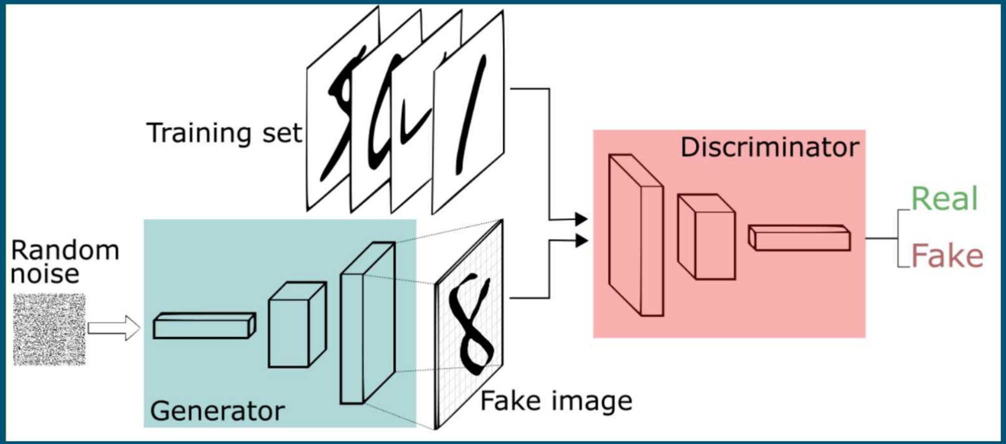

GAN diagram. The random noise refers to a latent variable

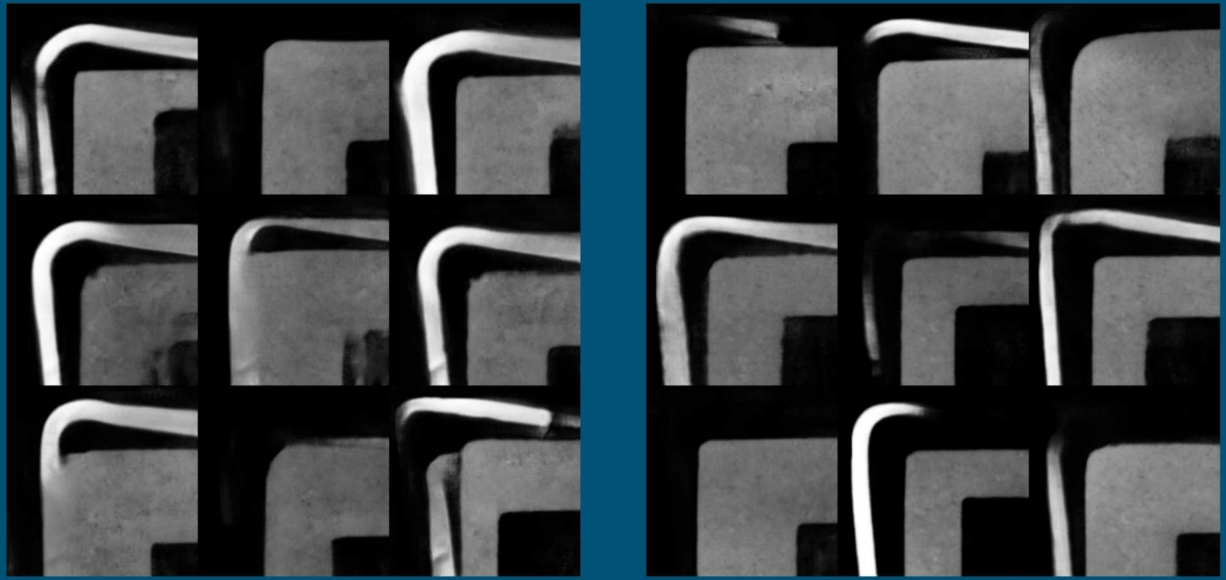

Reconstructions of trained AnoGAN with different latent dimensions. Left: 10-d latent space, Right: 30-d latent space

Structural Similarity improves with increasing latent dimension

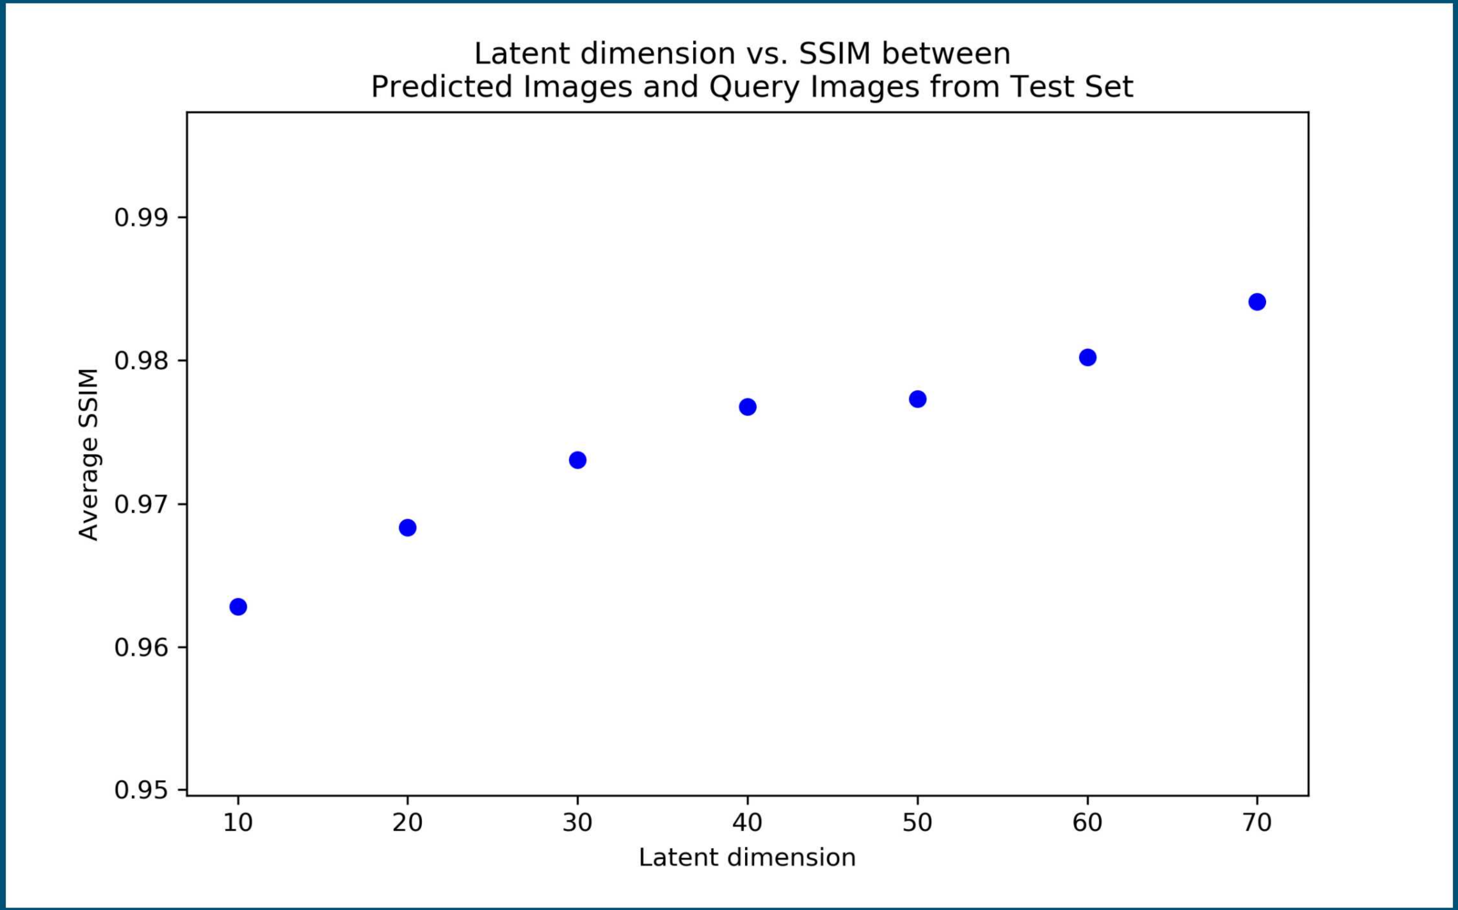

A larger latent space allows the generator to create a more complex and expressive mapping from latent vectors to the space of images

An improved metric is needed to distinguish anomalous from non-anomalous images

Anomaly score from original AnoGAN paper:

- $A\left(x_q, G(z_\Gamma)\right) = (1-\lambda)(\sum|x_q - G(z_\Gamma)|) + \lambda(\sum|f(x_q) - f(G(z_\Gamma))|)$
- $\sum|x_q - G(z_\Gamma)|$: Residual loss, measures absolute image differences
- $\sum|f(x_q) - f(G(z_\Gamma))|$: Discriminator loss, measures difference in discriminator layers

AnoGAN paper achieves 0.89 AUC from ROC between false positive rate and true positive rate

An improved metric is needed to distinguish anomalous from non-anomalous images

Anomaly score from original AnoGAN paper:

- $A\left(x_q, G(z_\Gamma)\right) = (1-\lambda)(\sum |x_q - G(z_\Gamma)|) + \lambda(\sum |f(x_q) - f(G(z_\Gamma))|)$
- $\sum |x_q - G(z_\Gamma)|$: Residual loss, measures absolute image differences
- $\sum |f(x_q) - f(G(z_\Gamma))|$: Discriminator loss, measures difference in discriminator layers

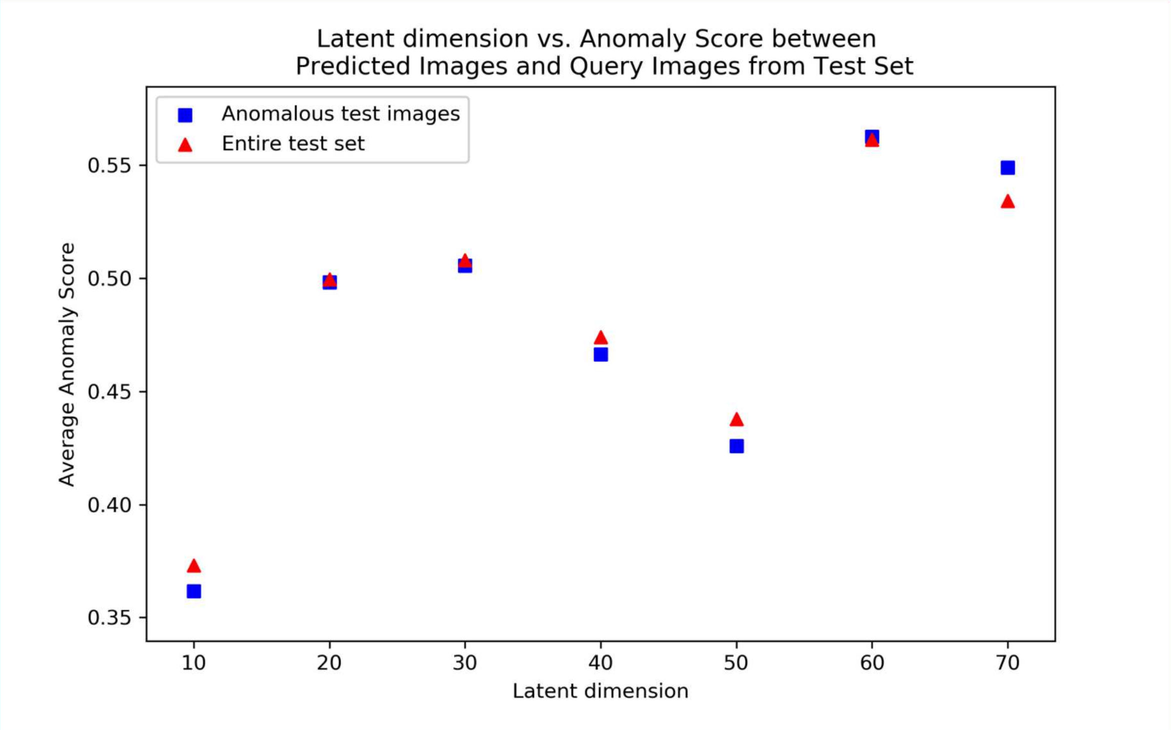

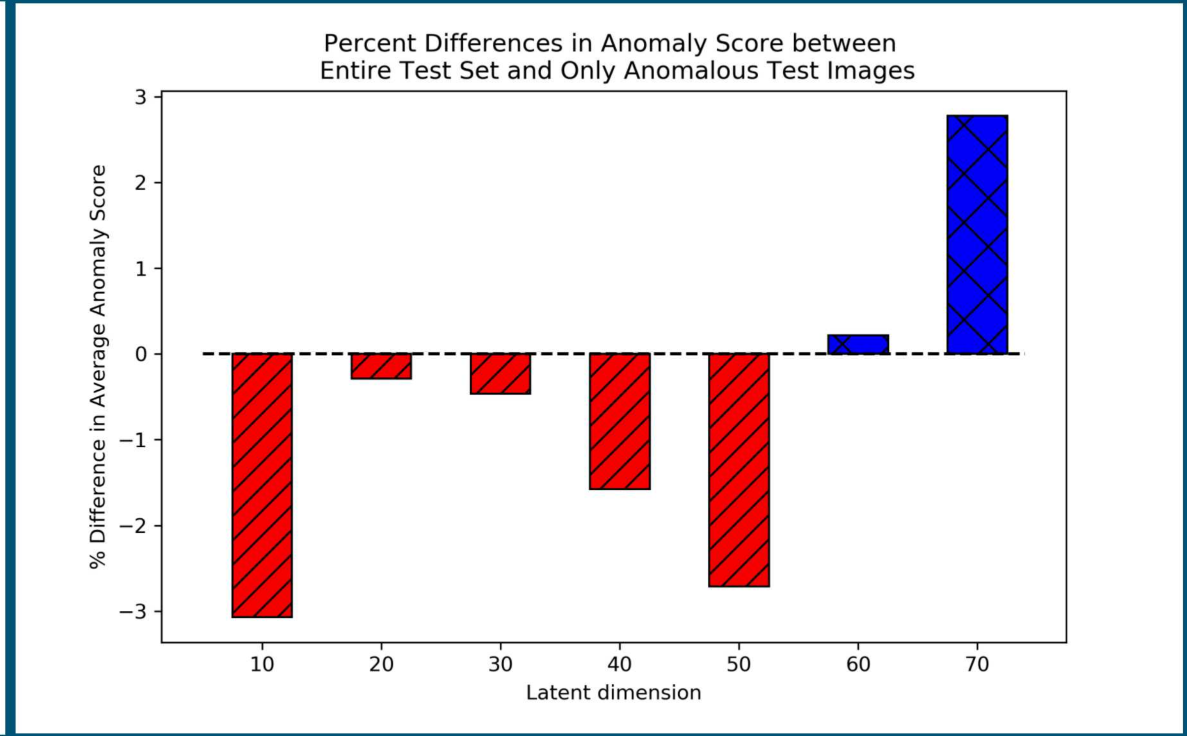

Many potential research directions could reach a solution

Improve generative model training

- Often, training runs suffered from mode collapse, resulting in poor images

Use an architecture with inverse latent vector to image mapping

- Avoid costly backpropagation at inference
- As in Schlegl et. al. "f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks"

Improved metric for distinguishing anomalous images from normal ones

- Investigate effect of tuning λ in anomaly score function
- Invent a new function for discovering fine-grained anomalies
- Use uncertainty quantification techniques for qualifying generated images that may be unrealistic

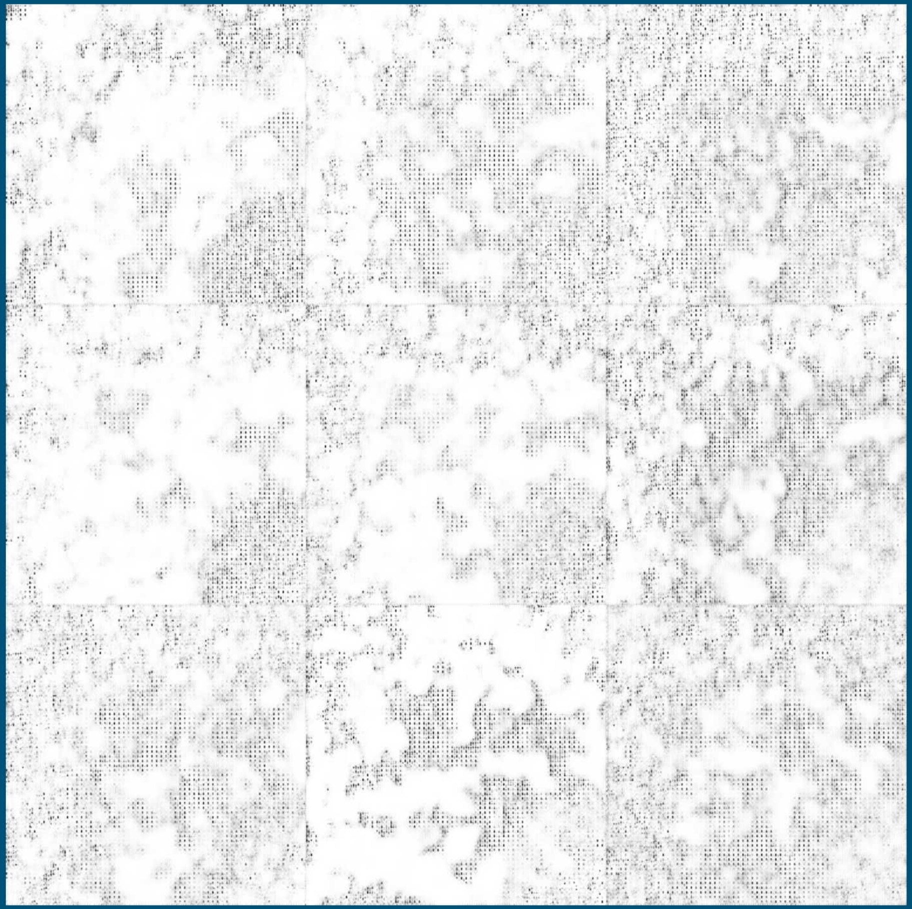

Thank you!

Electronics support and expertise:

- Christian Turner
- Mike Baker
- Norman Padilla

Sandia Interdisciplinary Machine Learning Research team:

- Kevin Potter
- Cari Martinez
- Matthew Smith