Deep Representation Learning
FY16-FY19

Jason E. Summers, ARiA

Objectives

Problem:

- Automatic target recognition (ATR) from synthetic aperture sonar (SAS) images for mine countermeasures (MCM) works well on small sets of labeled or numerically simulated data, but underperforms in operational environments

Solution:

- use deep-learning to find feature representations that account for context in measured data
- synthesize realistic data with learned generative representations to capture variability and diversity

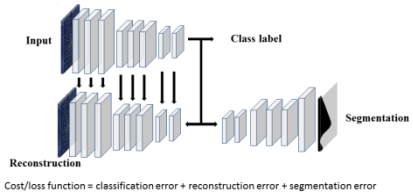

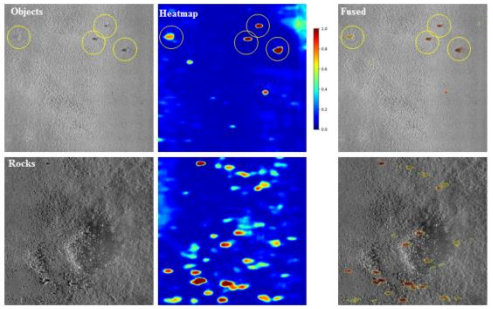

fully convolutional ladder network for pixelwise segmentation of SAS images using semisupervised training from sparsely labeled data

Technical Approach

Novel first uses in the MCM domain of emerging techniques

- use convolutional autoencoders to learn deep feature representations from sparsely labeled operational SAS data (FY16)
- develop fully convolutional ladder networks to enable semisupervised pixel-wise multiclass segmentation (FY17)
- understand output of deep networks using class-activation mapping (FY17)
- use generative adversarial networks (GANs) to learn deep generative representations from operational SAS data (FY16)
- match latent representations in deep convolutional networks to measured and simulated images to synthesize images and transfer style from measured data to numerically simulated data (FY16)
- learn physically interpretable features with InfoGANs (FY18)

Accomplishments

FY16:

- CAEs learn interpretable features from operational data that discriminate bottom types
- GANs generate simulations from learned representation that image classifier cannot distinguish from real
- convolutional networks synthesize images using style learned from single SAS images and transfer style to simulated targets

FY17:

- Novel fully convolutional ladder networks for semisupervised pixel-wise multi-class segmentation of whole images in on feed-forward pass
- Class-activation mapping for explaining network decisions

September 2017