Laboratories

SAND2019-9435C



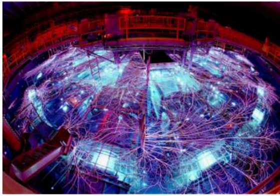

Deep neural networks for compressive hyperspectral imaging

Dennis J. Lee

SPIE 2019

Compressive hyperspectral imaging reduces measurements

- Compressive sensing reduces the number of bands from the full spectra.
- Less measurements results in faster acquisition time.
- Reconstruction requires heavy computation.
- Can neural networks reduce computation time for reconstruction?
- Can neural networks improve classification accuracy of spectra?
- Can neural networks classify the raw, compressed spectra?

Outline: How to reconstruct and classify compressive hyperspectral images?

1. Hyperspectral imaging based on compressive sensing
2. Task 1: Reconstruction of hyperspectral images from compressive sensing data
3. Task 2: Classification of hyperspectral images from compressive sensing data

Hyperspectral imaging based on compressive sensing

Hyperspectral imager based on compressive sensing

- Measure the spectra of incoming light using a spatial light modulator (SLM) before the focal plane array.
- The transmission of the SLM can be controlled by applying different voltages.
- Each measurement corresponds to a different transmission profile of the SLM.

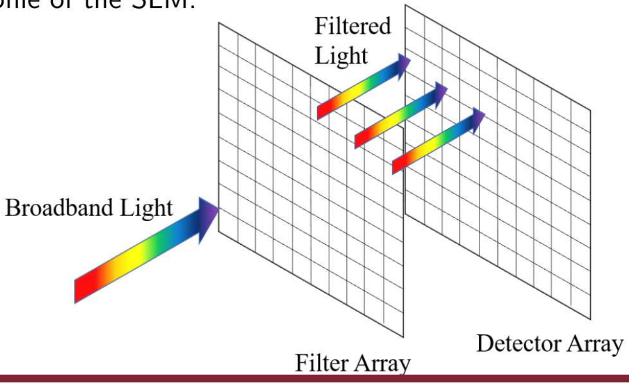

Simulate transmission of Fabry Perot resonators

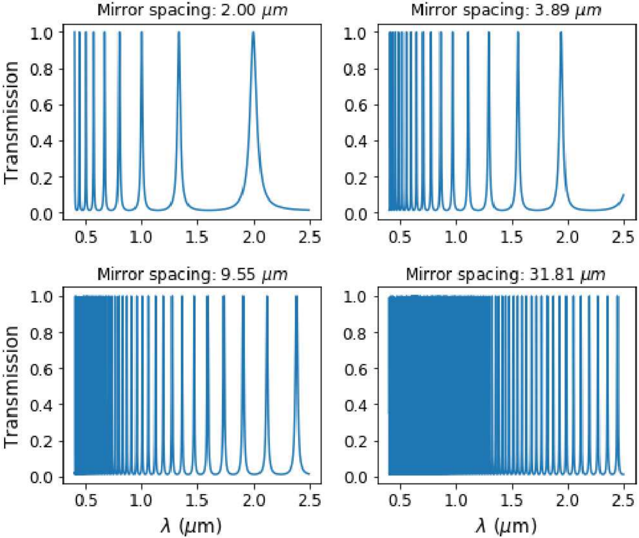

D. J. Lee and E. A. Shields, "Compressive hyperspectral imaging using total variation minimization" SPIE (2018).

Compress measurements using Fabry Perot resonators

- Each measurement corresponds to a different transmission profile.
- Take less measurements than the number of bands.

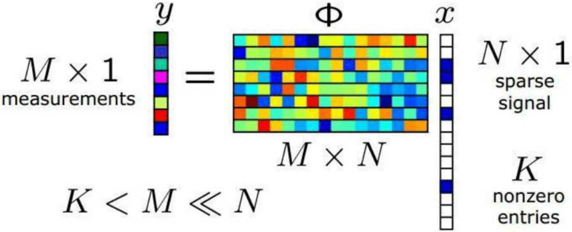

Sample mirror spacing of Fabry Perot resonators

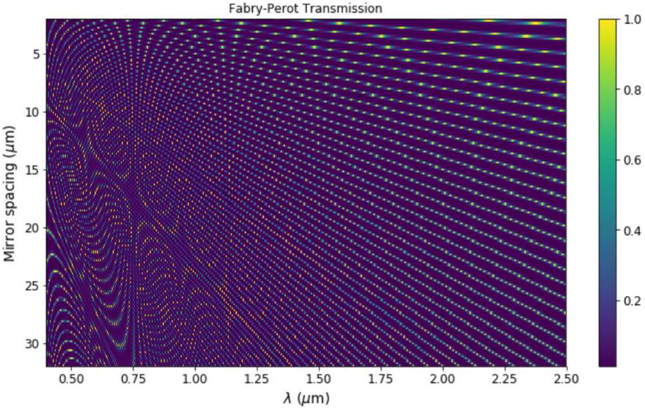

Task 1: Reconstruction of hyperspectral images from compressive sensing data

Multi-layer perceptrons reconstruct hyperspectral images

- The output is the full hyperspectral image with 220 bands.
- The inputs are the compressed hyperspectral images with the number of bands varying from 160, 80, 40, 20, and 10.
- Vary the number of layers: K = 1, 2, 4, 7, 14.

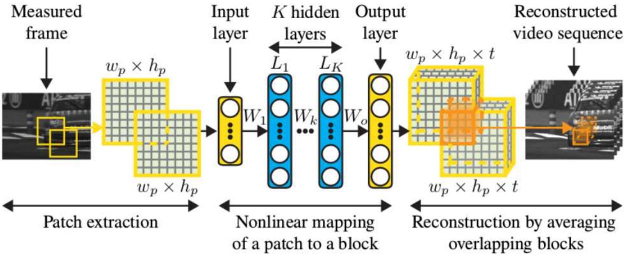

M. Iliadis, et. al., "Deep fully-connected networks for video compressive sensing" arxiv:1603.04930v2 (2017).

The Indian Pines dataset is a hyperspectral image

- Dataset consists of 145 x 145 pixels with a spatial resolution of 20 m and a 10 nm spectral resolution over the range of 400–2500 nm.
- We use the entire 220 bands including the water absorption region. This simulates the real application, where the entire spectrum is modulated by a LCD.

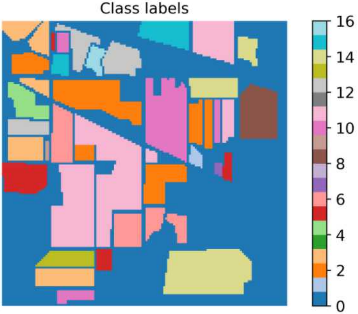

Train on 2 datasets: Indian Pines and random spectra

- Indian Pines contains 21025 spectra (145 × 145). We reserve 60% for training, 20% for validation, and 20% for testing.
- We generate random spectra from a normal distribution (unit mean, zero variance), smoothed by a Hanning filter with window sizes that vary from 11, 21, 31, and 41. There are equal numbers of Indian Pines and random spectra.
- Random noise is added to the Indian Pines training data, generated from the distribution described above.
- The training set is split 50% between Indian Pines and random spectra.
- The random dataset and additive noise help to prevent overfitting. Both datasets are normalized to zero mean and unit variance.

Moderate compression (160/220) => small error

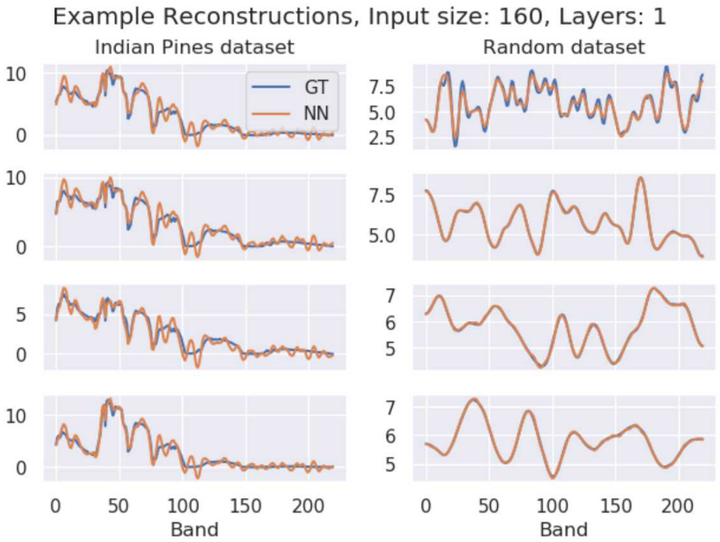

Larger compression (10/220) => larger error

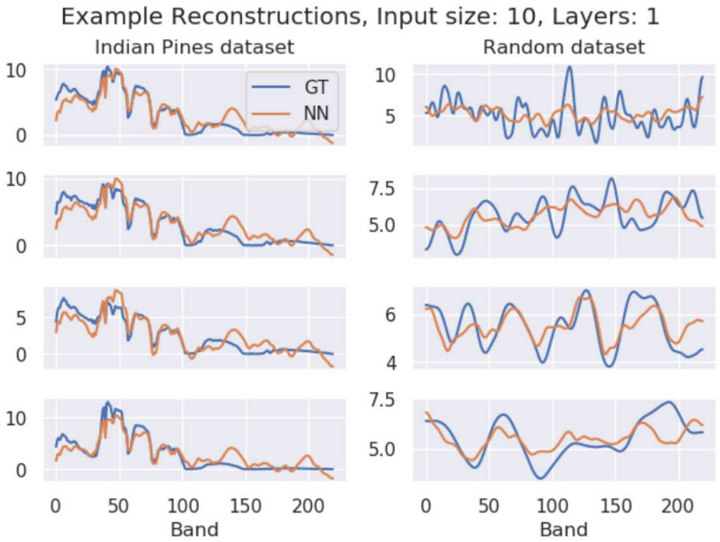

Crop spectra vary with input size of 10 and single layer

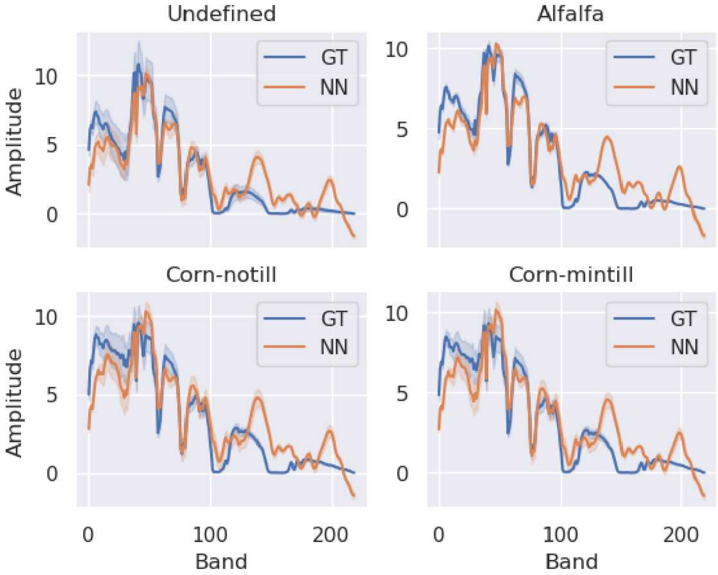

Crops with fewer examples show higher variance

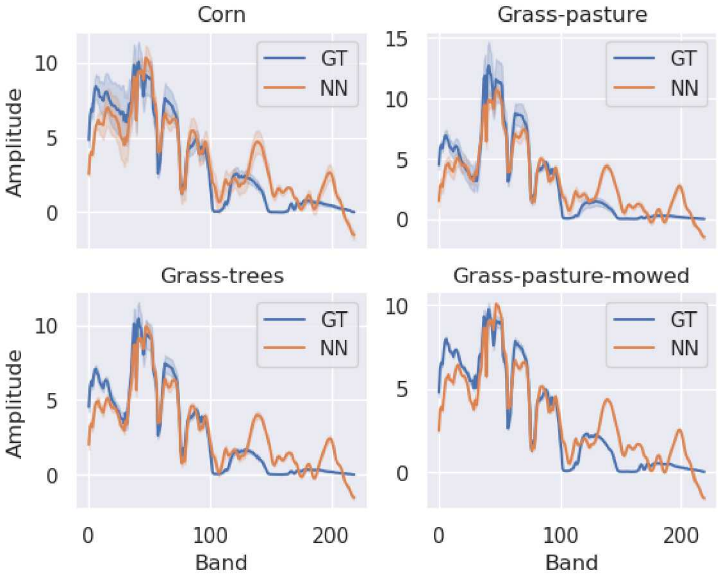

Higher bands have lower signal, greater error

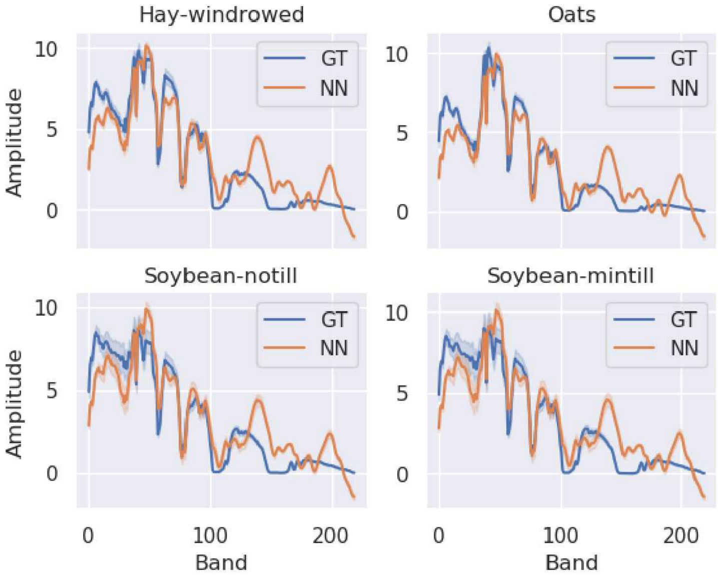

Stone-steel-towers has fewest examples, greatest error

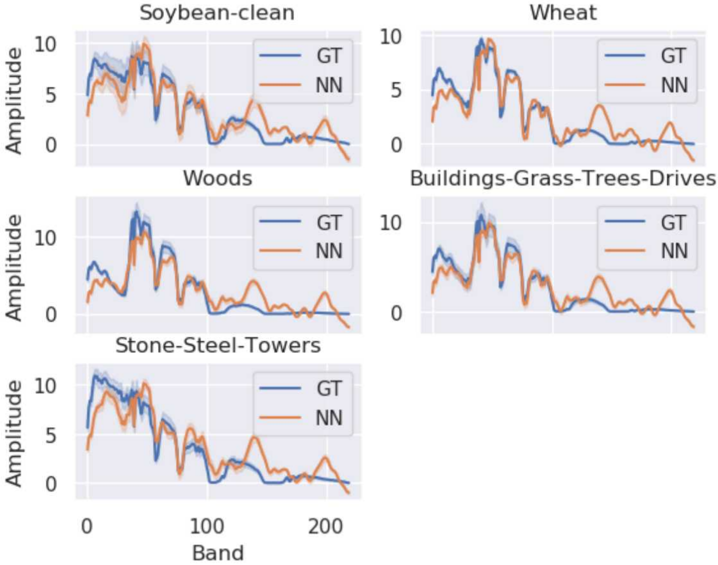

Overall R^2 shows too many layers increases error

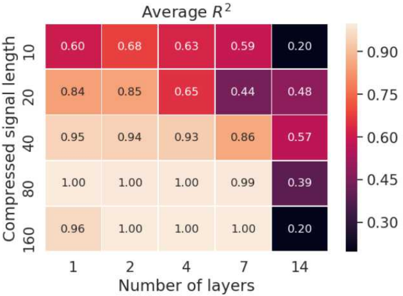

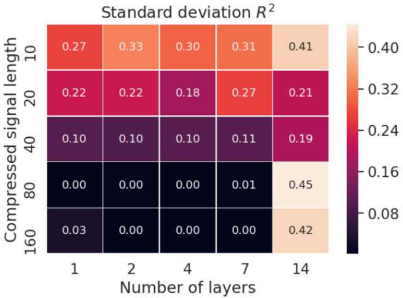

Single layer shows least overfitting

Regularize the model to further reduce overfitting

Sandia National Laboratories

- With a single layer, the Indian Pines reconstruction seemed to show the least overfitting.
- The random dataset serves as a measure of overfitting.
- Can the training dataset be further augmented beyond random spectra?
- How can the multilayer perceptron be further regularized? Some ideas are to add drpout, noise augmentation.

Sandia National Laboratories

Task 2: Classification of hyperspectral images from compressive sensing measurements

Evaluate classifier performance on compressed inputs

- Compress the full spectra (220 bands) to 160, 80, 40, 20, and 10 bands.
- Apply a classifier (SVM, KNN, Neural networks) to the reconstructed spectra (220 bands).
- Evaluate classifier performance on compressed spectra (160, 80, 40, 20, 10 bands)..
- Evaluate reconstructed spectra (220 bands) as initial compressed input varies (160, 80, 40, 20, 10 bands).

3D convolutions extract spatial and spectral features

- Two convolutional layers have kernels of size 3 x 3 x 7 and 3 x 3 x 3.
- Previous approaches have applied principle components analysis to the spectral dimension independently of the spatial dimension.

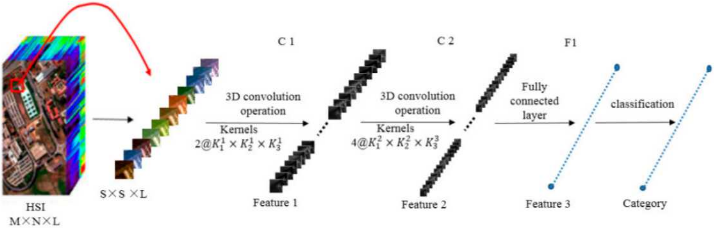

Y. Li, et. al., "Spectral-spatial classification of hyperspectral imagery with 3D convolutional neural network" Remote Sens. (2017).

Extract multiscale spatial and spectral features

- Evaluate a multi-scale 3D deep convolutional neural network on compressed vs. reconstructed spectra.
- Convolutions occur in 3D across the spatial and spectral dimensions.
- The kernel size varies along the spectral dimension (1, 3, 5, 11).

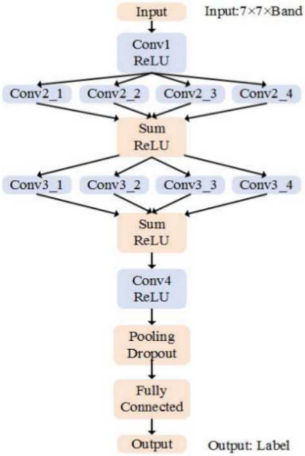

M. He, et. al., "Multi-scale 3D deep convolutional neural network for hyperspectral image classification" ICIP (2017).

Recurrent networks characterize spectral correlation

- Gated recurrent units predict bands of each hyperspectral pixel.
- Model accounts for spectral but not spatial correlations.

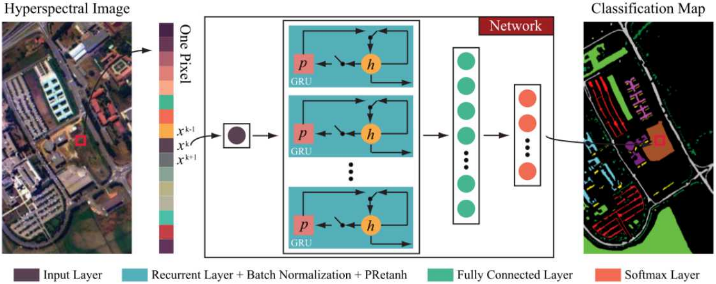

L. Mou, et. al., "Deep recurrent neural networks for hyperspectral image classification" IEEE Trans. Geoscience and Remote Sens. (2017).

Recurrent networks performs best on compressed input

- Metrics in paper: overall accuracy, average accuracy, kappa.
- Compressed spectra may lose spatial context.
- For SVM on compressed inputs, we use a RBF kernel with C $= 1000$, $\gamma = 0.001$, determined from a grid search.
- For KNN, we search over K $=$ 1, 3, 5, 10, 20. In most cases, K $=$ 10 performed best (eg, for compressed sizes of 10 and 160).

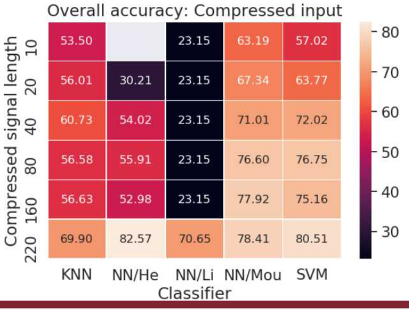

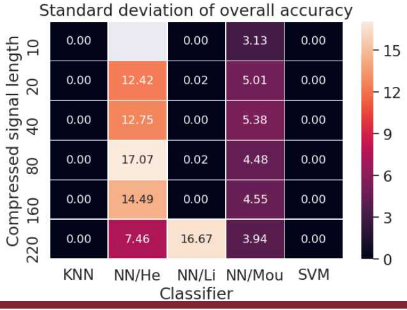

How do RNNs perform on compressed inputs?

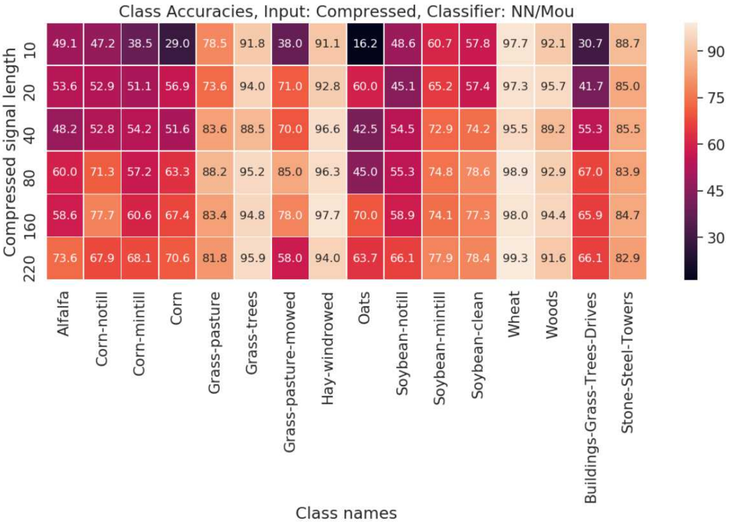

3D CNNs perform best on the reconstructed inputs

- Metrics in paper: overall accuracy, average accuracy, kappa.
- Reconstructed image contains meaningful spatial information.
- For SVM on compressed inputs, we use a RBF kernel with C $= 100$, $\gamma = 0.01$, determined from a grid search.
- For KNN, we search over K $=$ 1, 3, 5, 10, 20. Usually, K $=$ 5, 10 performed best (eg, for sizes 160, 10 respectively).

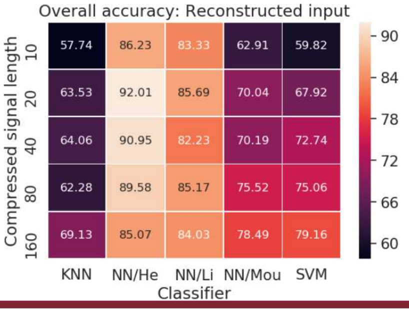

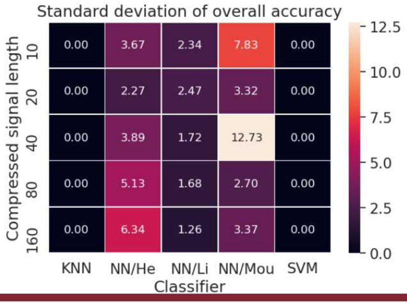

How do 3D CNNs perform on reconstructed inputs?

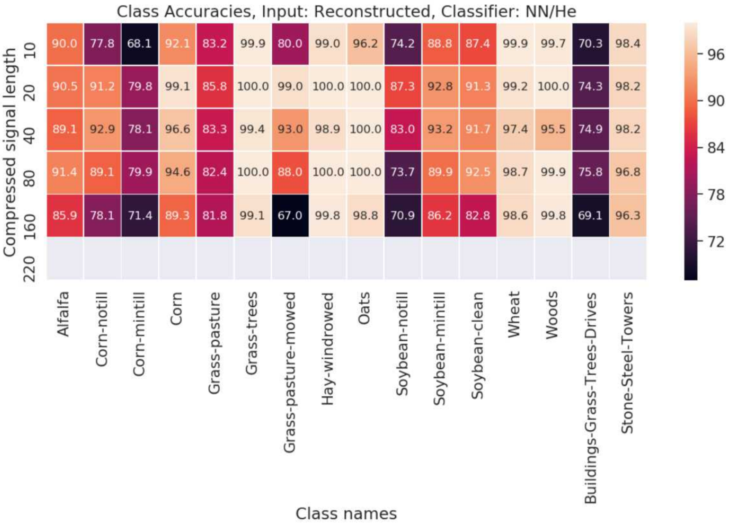

Accuracy improves on reconstructed inputs

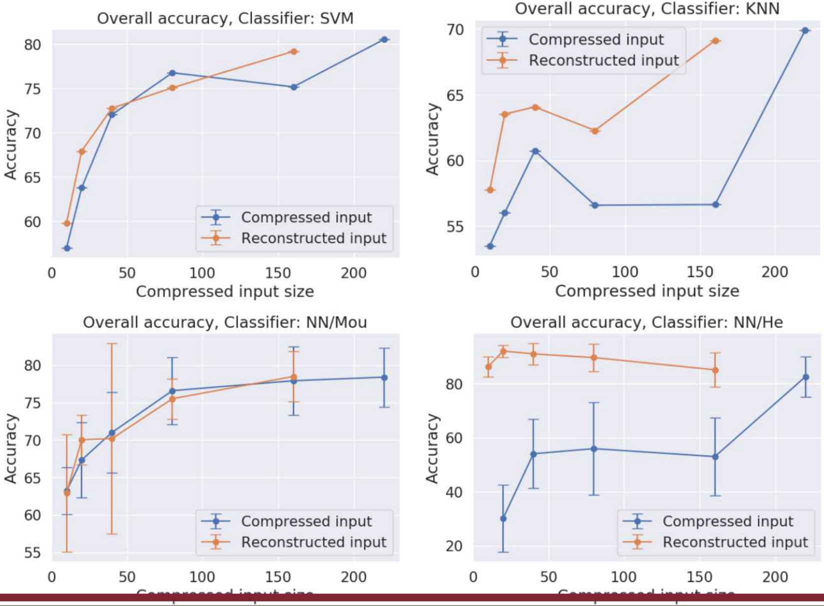

Why does accuracy improve on reconstructed spectra?

Sandia National Laboratories

- The multilayer perceptron reconstructs spectra accurately even for high compression.
- The compressed input may lost spatial context, so 3D convolutions may not be as effective.
- Compressive sensing can be an unstable inverse problem. Similar spectra can look completely different in compressed space.

Summary: Neural networks can reconstruct and classify compressive hyperspectral images

- We have demonstrated a two step process for hyperspectral image classification using compressive sensing measurements.
- First reconstruct the hyperspectral image from compressive sensing measurements. We investigated varying numbers of layers in a multiayer perceptron. We found that 1 layer minimizes overfitting.
- Second classify the reconstructed image. We explored SVMs, KNN, and 3 neural networks (3D CNNs, Multiscale 3D CNNs, RNNs). We calculated accuracy using the compressed (non-reconstructed) measurements. We found that accuracy improves using reconstructed spectra compared to raw compressed measurements.