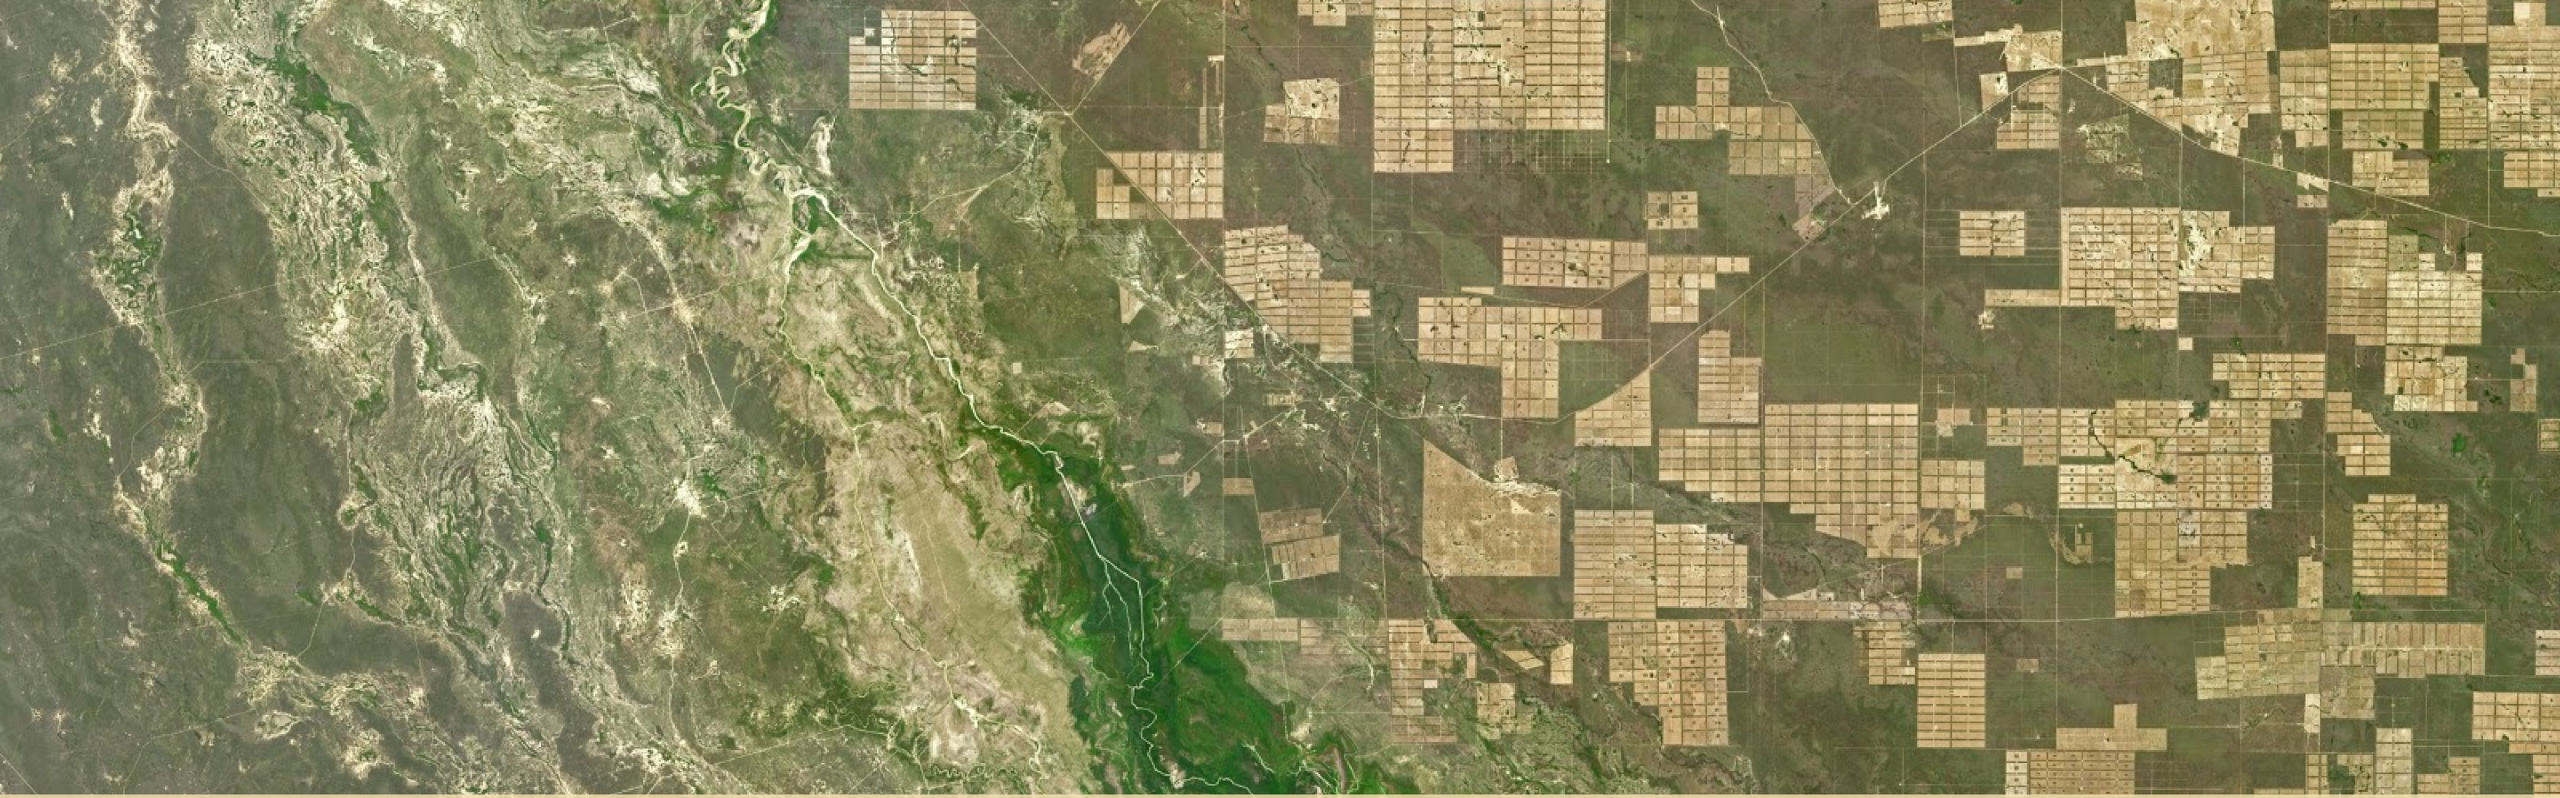

Machine Learning for Remote Sensing Applications

Hannah R. Kerner
Asst Research Professor, University of Maryland – College Park
COSG Fall Meeting 2019

HARVEST
EARTH DATA
FOR INFORMED AGRICULTURAL DECISIONS

Remote sensing: perfect storm for ML

- Huge data volumes
  - 11 TB/day from Planet Labs alone

San Francisco, February 11, 2017. Planet, Inc.

Remote sensing: perfect storm for ML

- Huge data volumes
  - 11 TB/day from Planet Labs alone
- High-dimensional data
  - Multispectral, hyperspectral

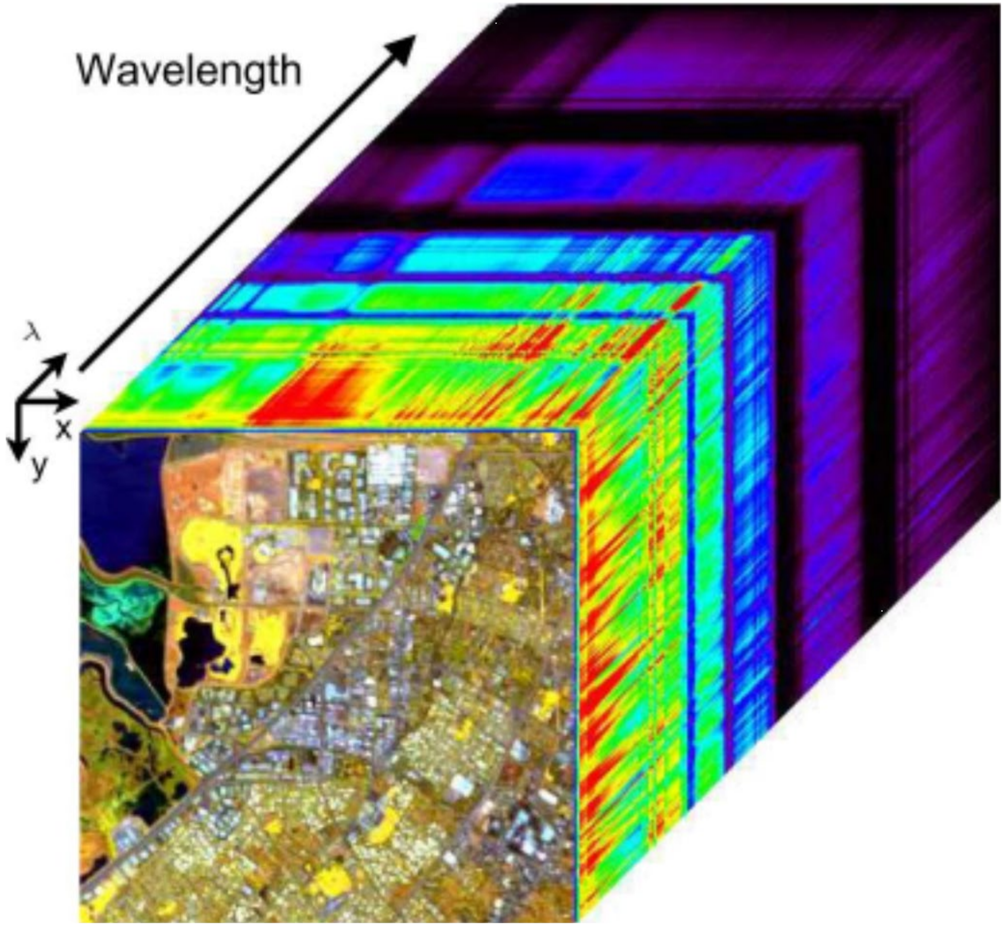

Credit: Christophe et al., 2008

Remote sensing: perfect storm for ML

- Huge data volumes
  - 11 TB/day from Planet Labs alone
- High-dimensional data
  - Multispectral, hyperspectral
  - Frequent revisit times
    - Landsat: 16 days
    - Sentinel-2: 5 days
    - Planet Labs: daily or sub-daily

Viedma Glacier, Southern Patagonia, South America. Credit: Planet, Inc.

Remote sensing: perfect storm for ML

- Huge data volumes
  - 11 TB/day from Planet Labs alone
- High-dimensional data
  - Multispectral, hyperspectral
  - Frequent revisit times
    - Landsat: 16 days
    - Sentinel-2: 5 days
    - Planet Labs: daily or sub-daily
- Complex relationships within and between spatial, spectral, temporal dimensions

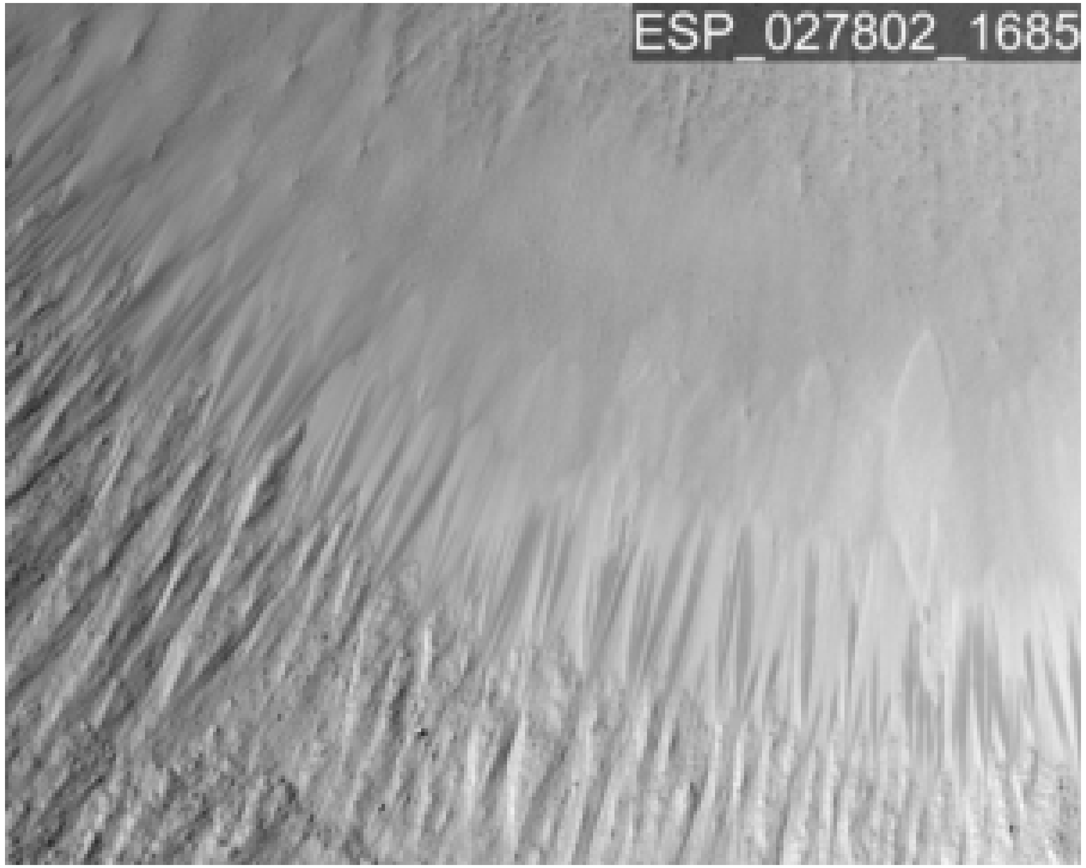

Credit: UA/LPL/NASA/Brian Bue

Challenges for ML in remote sensing

- Datasets not independently and identically distributed (i.i.d.)
    - Class imbalances
    - Difficult to assess generalization since objects represented by multiple pixels

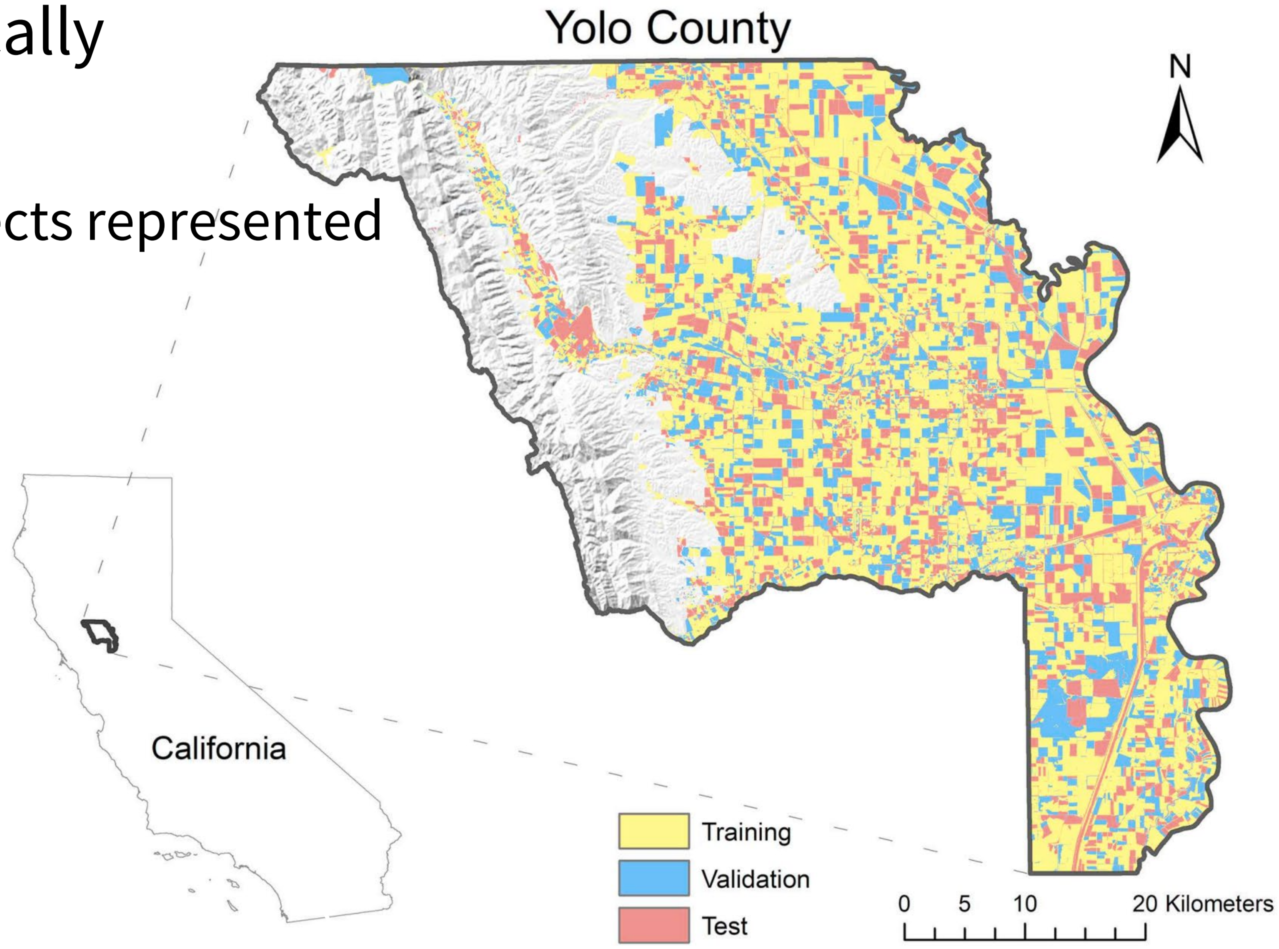

Zhong et al., 2019

Challenges for ML in remote sensing

- Datasets not independently and identically distributed (i.i.d.)
  - Class imbalances
  - Difficult to assess generalization since objects represented by multiple pixels
- Tiles much larger than typical ML image sizes

CIFAR10 images: 32×32×3 px

Landsat-8 image: 3660×3660×11 px

(Harmonized Landsat and Sentinel-2 image over N Illinois. Credit: USGS/NASA

Challenges for ML in remote sensing

- Datasets not independently and identically distributed (i.i.d.)
  - Class imbalances
  - Difficult to assess generalization since objects represented by multiple pixels
- Tiles much larger than typical ML image sizes
- Non-trivial pre-processing/cleaning
  - Cloud removal
  - Orbital track
  - Interpolation/smoothing
  - Co-registration

Sentinel-2 Time Lapse (Cropped)
Harmonized Landsat and Sentinel-2 (HLS)
Credit: USGS/NASA

Challenges for ML in remote sensing

- Datasets not independently and identically distributed (i.i.d.)
  - Class imbalances
  - Difficult to assess generalization since objects represente by multiple pixels
- Tiles much larger than typical ML image sizes
- Non-trivial pre-processing/cleaning
  - Cloud removal
  - Orbital track
  - Interpolation/smoothing
  - Co-registration
- Labeling often requires domain expertise

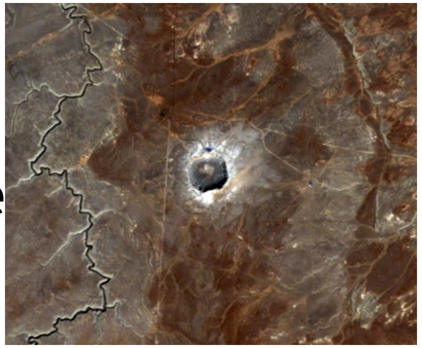

Meteor (Impact) Crater, AZ.
(Landsat, USGS/NASA)

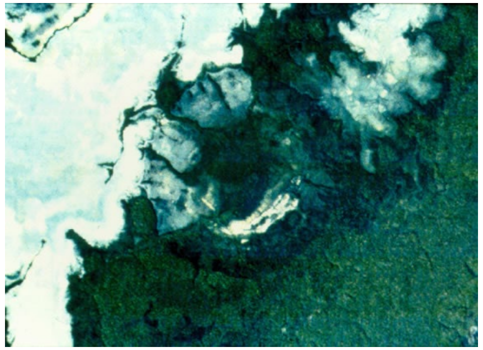

Iturralde (Suspected) Crater, Bolivia.
(Landsat, USGS/NASA)

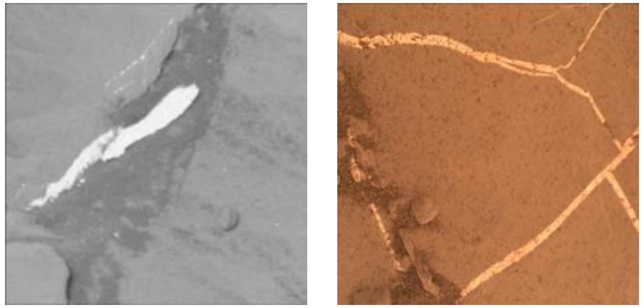

Veins or bright-toned material?
Mars Science Laboratory, NASA/JPL

Challenges for ML in remote sensing

- Datasets not independently and identically distributed (i.i.d.)
  - Class imbalances
  - Difficult to assess generalization since objects represent by multiple pixels
- Tiles much larger than typical ML image sizes
- Non-trivial pre-processing/cleaning
  - Cloud removal
  - Orbital track
  - Interpolation/smoothing
  - Co-registration
- Labeling often requires domain expertise
- Input sources have different physical units

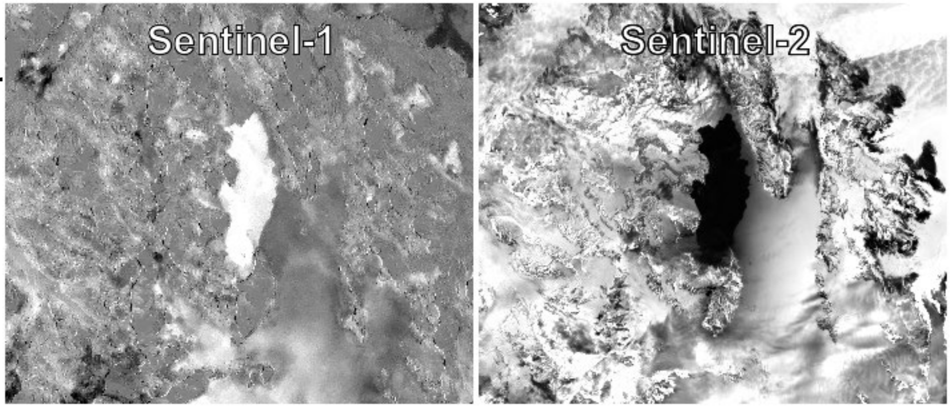

Sentinel-1 (radar) and Sentinel-2 (multispectral) images of Glacier Bay landslide
Credit: ESA/Simon Gascoin

Challenges for ML in remote sensing

- Datasets not independently and identically distributed (i.i.d.)
  - Class imbalances
  - Difficult to assess generalization since objects represented by multiple pixels
- Tiles much larger than typical ML image sizes
- Non-trivial pre-processing/cleaning
  - Cloud removal
  - Orbital track
  - Interpolation/smoothing
  - Co-registration
- Labeling often requires domain expertise
- Input sources have different physical units
- Differences in illumination conditions

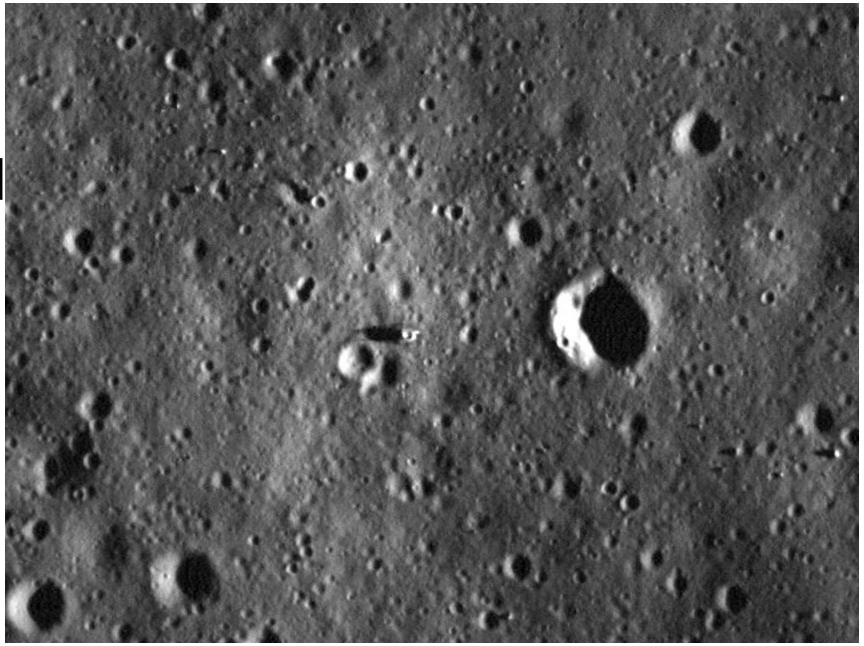

LROC images of Apollo 11 landing site from lunar dusk to dawn
Credit: NASA/ASU/Rob Pettengill

Remote sensing applications of ML

Common applications

- Land use/cover classification
- Scene classification
- Change detection/monitoring
- Anomaly/novelty detection
- Estimation of physical quantities

Emerging applications

- Object detection/mapping
- (Semantic) segmentation
- Pansharpening/super-resolution
- Registration

Land use and land cover classification

- Multi-class classification at pixel level
- Inputs typically vector or patch

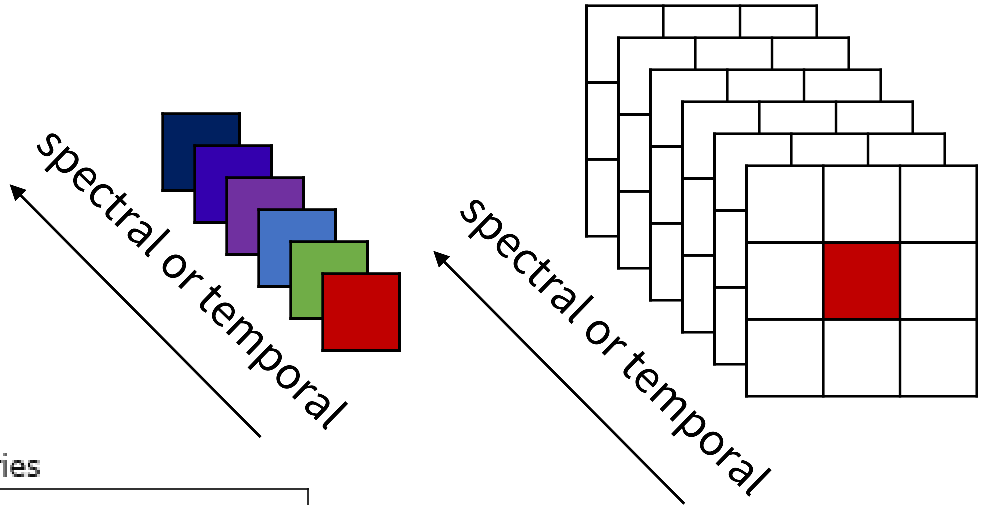

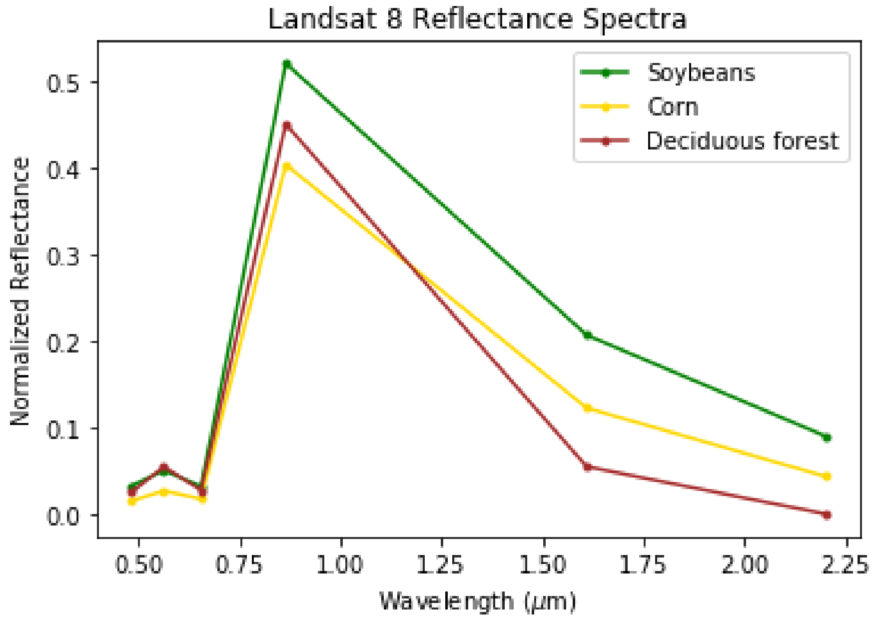

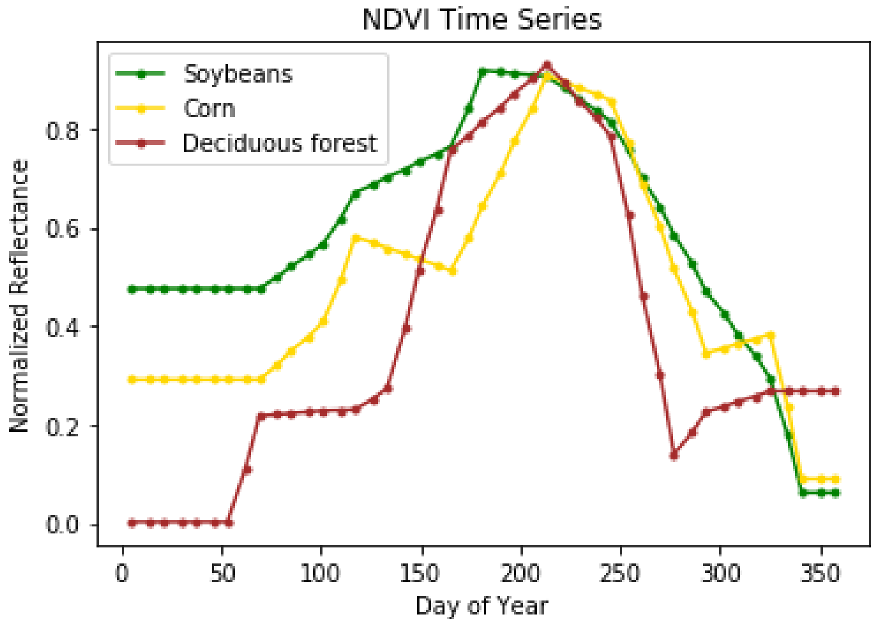

Land use and land cover classification

- Multi-class classification at pixel level
- Inputs typically vector or patch
- Labels from national databases (e.g., NLCD, CDL) or field campaigns
- Decision trees and random forests most common (supervised)
- Deep learning gaining popularity
  - Convolutional neural networks (1D/2D)
  - Recurrent neural networks/LSTMs

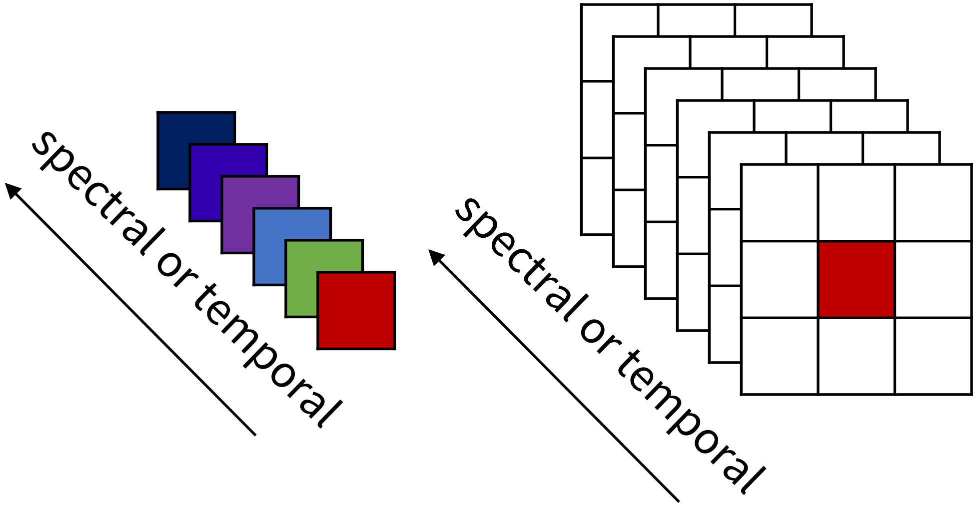

Target: reference data (ground-truth)

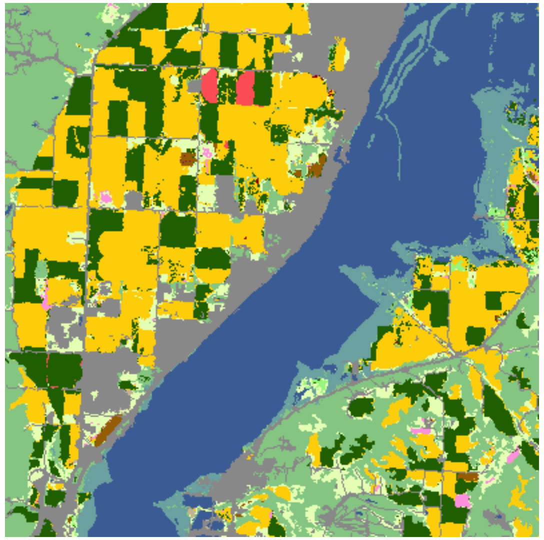

USDA/NASS Cropland Data Layer (N Illinois, 2017)

Land use and land cover classification

- Multi-class classification at pixel level
- Inputs typically vector or patch
- Labels from national databases (e.g., NLCD, CDL) or field campaigns
- Decision trees and random forests most common (supervised)
- Deep learning gaining popularity
  - Convolutional neural networks (1D/2D)
  - Recurrent neural networks/LSTMs

Zhong et al., 2018
Methods: compared 1D CNN and LSTM

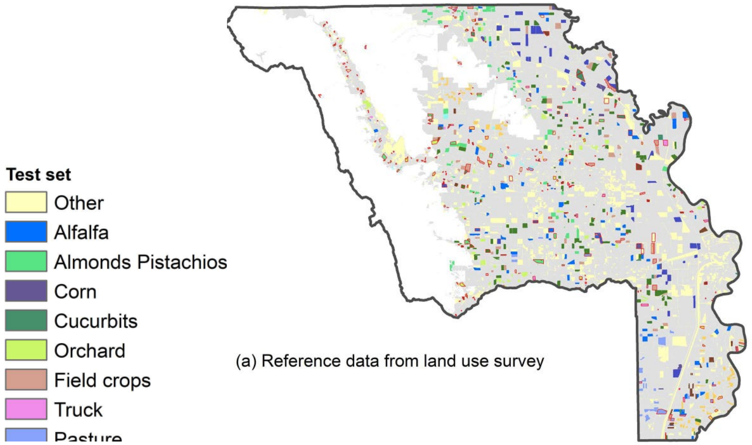

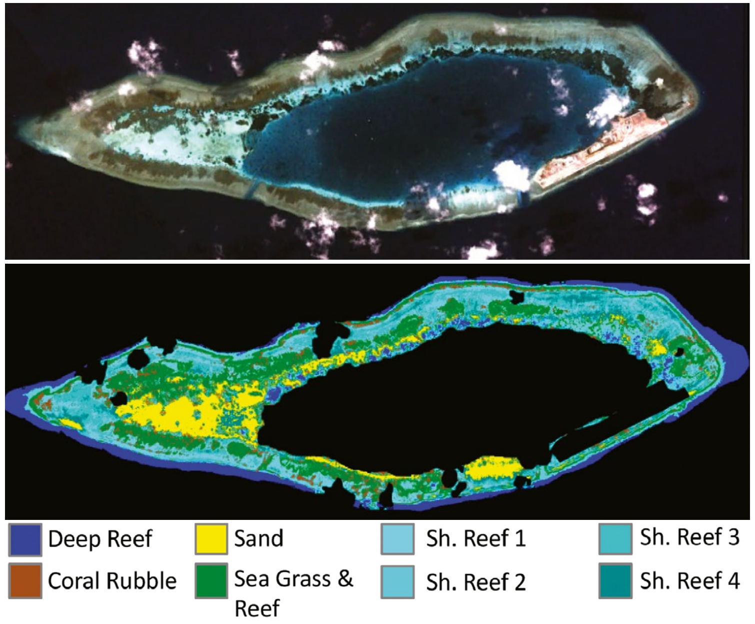

Asner et al., 2017
Method: K-means clustering (unsupervised)

Change detection and monitoring

- Pixel or image level
- Difference-based methods and post-classification comparison most common for pixel-level

Difference-based:

Before After Difference

Kerner et al., 2019

Post-classification comparison:

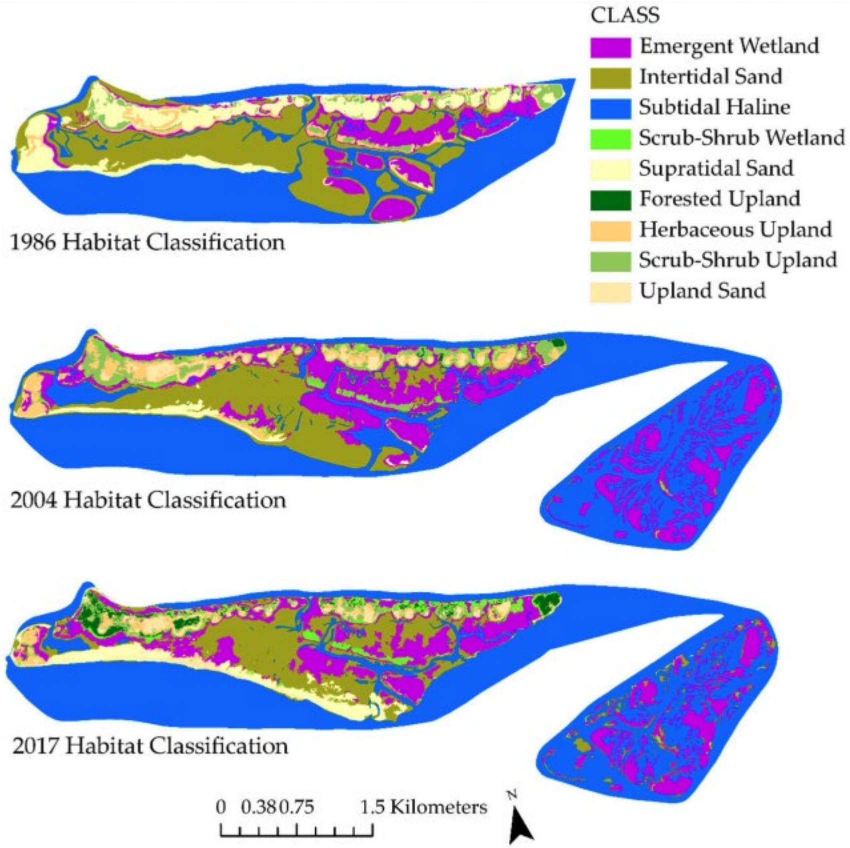

Gray et al., 2018

Change detection and monitoring

- Pixel or image level
- Difference-based methods and post-classification comparison most common for pixel-level
- Object-based and deep learning approaches for image-level
  - Multi-temporal images compared at feature level

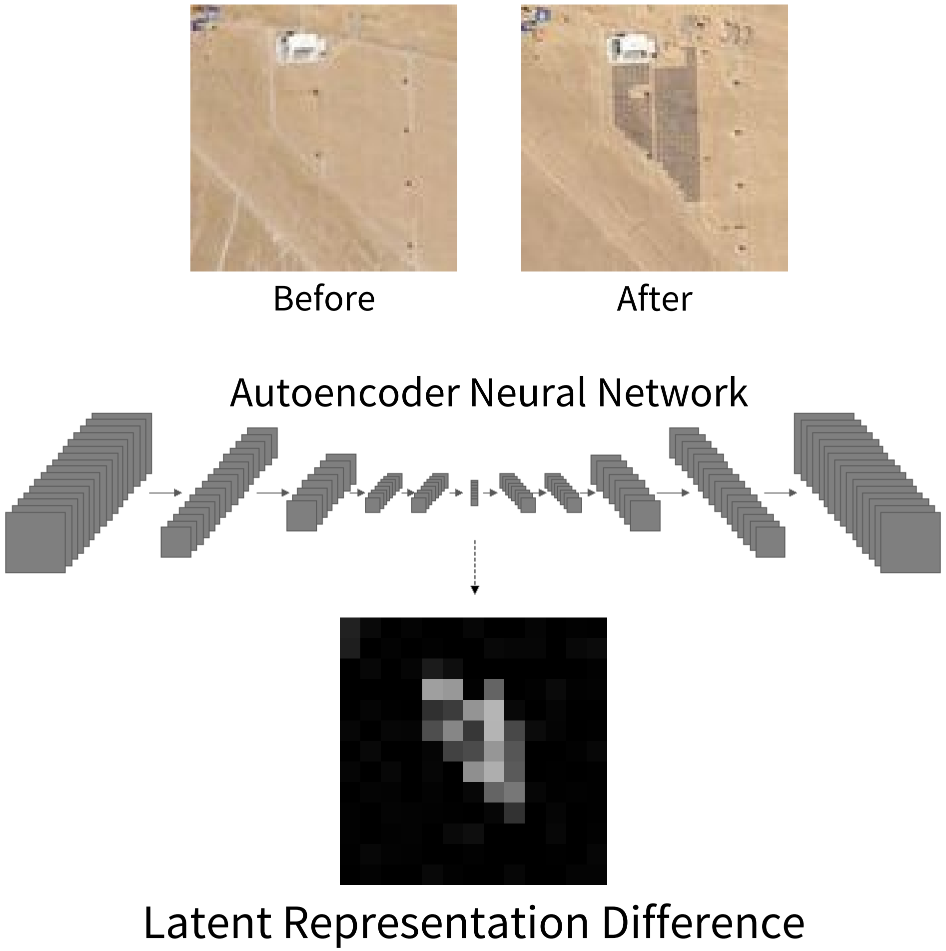

Kerner et al., 2019

Scene (image) classification

- Becoming more common with increasing spatial resolution
- Primarily enabled by deep learning, esp. convolutional neural networks
- Less labeled training data available
  - Some work on transfer learning using models pre-trained on ImageNet (14M)
    - e.g., Marmanis et al., 2015; Penatti et al., 2015; Castelluccio et al., 2015; Hu et al., 2015

Benchmark dataset collected from Google Earth:

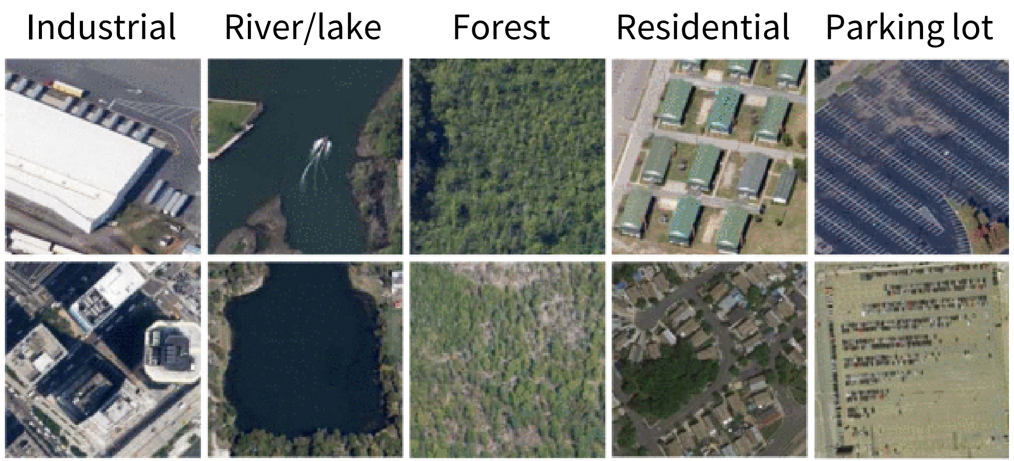

Zou et al., 2015

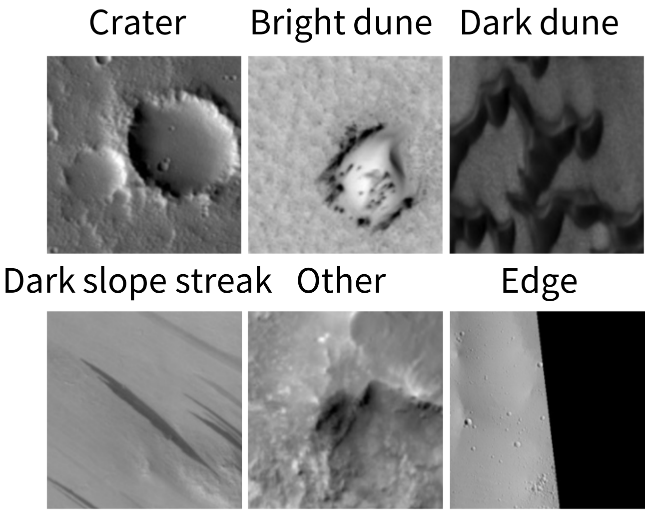

Scene classification used for Planetary Data System search

Wagstaff et al., 2018

Novelty/anomaly detection

- Detecting rare or unseen patterns (could be spatial, spectral, temporal)
- Typically unsupervised or one-class supervision (only “normal” known)
- Majority of remote sensing applications use Reed Xiaoli (RX) detector for pixel-wise anomaly scores

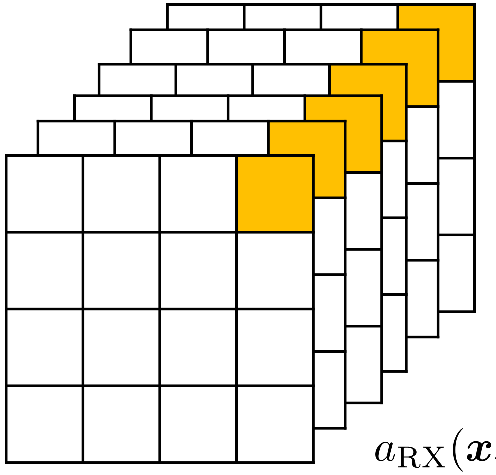

Mahalanobis distance between **pixel** and background distribution (Reed & Yu, 1990)

$$a_{\mathrm{RX}}(\boldsymbol{x}_i) = (\boldsymbol{x}_i - \boldsymbol{\mu}_t)^T \boldsymbol{\Sigma}_t^{-1} (\boldsymbol{x}_i - \boldsymbol{\mu}_t)$$

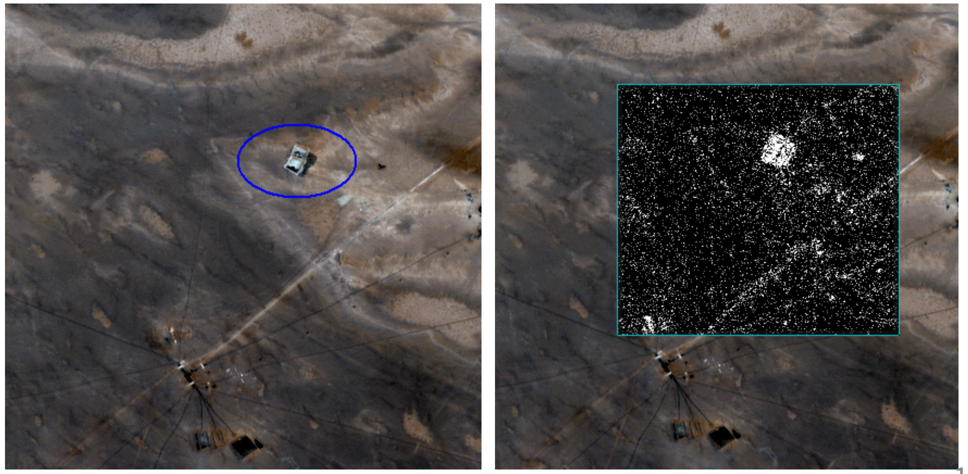

Credit: Harris Geospatial

Novelty/anomaly detection

- Detecting rare or unseen patterns (could be spatial, spectral, temporal)
- Typically unsupervised or one-class supervision (only “normal” known)
- Majority of remote sensing applications use Reed Xiaoli (RX) detector for pixel-wise anomaly scores
- Reconstruction-based methods common in ML literature

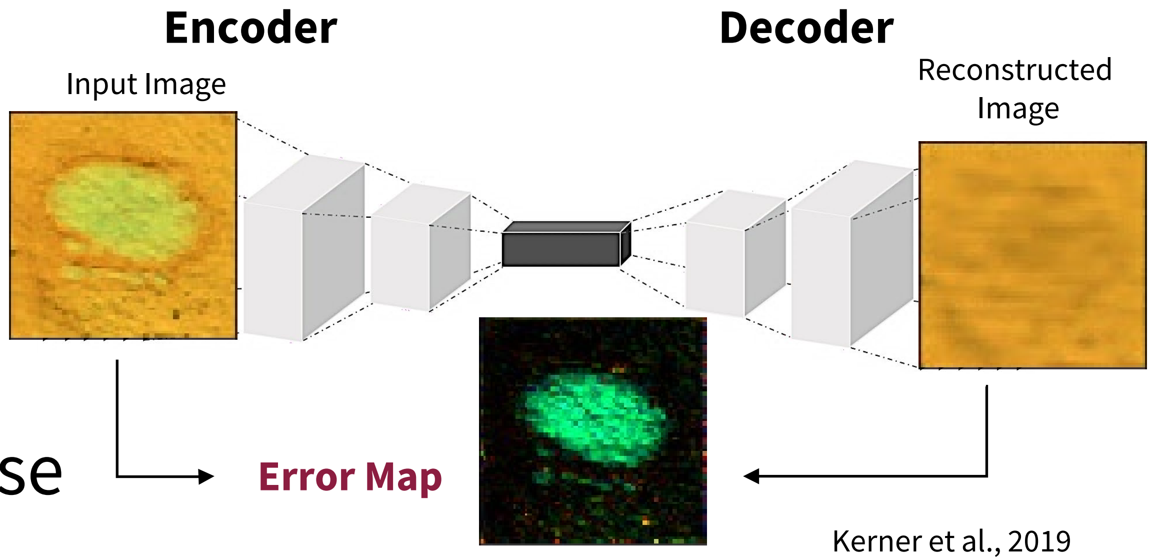

Kerner et al., 2019

Estimation of physical quantities

- Regression of physical/biophysical values directly from Earth observation data
- Often combines multiple data sources
- Common: random forests/regression trees, feed-forward neural networks, process models
- Emerging: convolutional neural networks, recurrent neural networks (LSTMs)

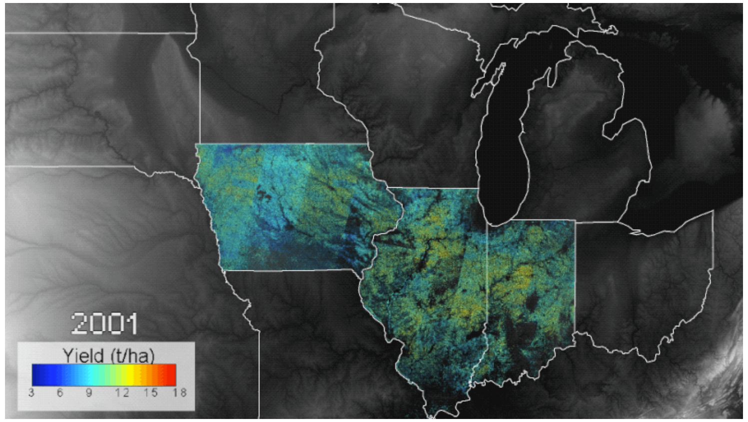

Maize yield estimates in US Corn Belt (Jin et al., 2017)

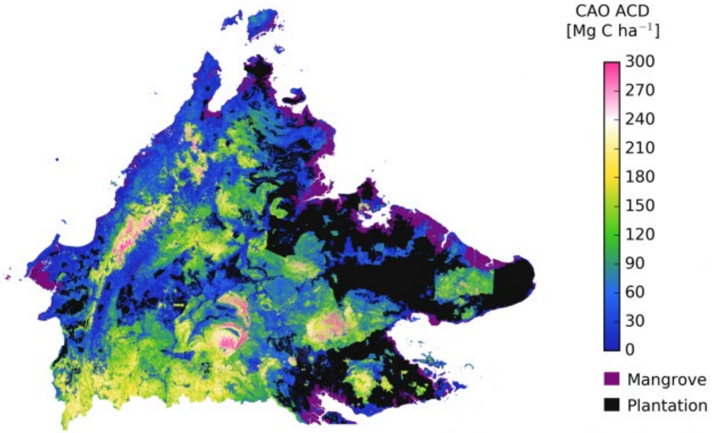

Aboveground carbon density in Borneo (Asner et al., 2018)

Remote sensing applications of ML

Common applications

- Land use/cover classification
- Scene classification
- Change detection/monitoring
- Anomaly/novelty detection
- Estimation of physical quantities

Emerging applications

- Object detection/mapping
- (Semantic) segmentation
- Pansharpening/super-resolution
- Registration

Object detection/mapping

- Bounding box predicted around feature or object of interest
- Common deep learning architectures (Zhao et al., 2019):
  - YOLO (v3)
  - R-CNN/Fast R-CNN/Faster R-CNN
  - Regional Fully Connected Networks (R-FCN)
- Challenge: redundancy

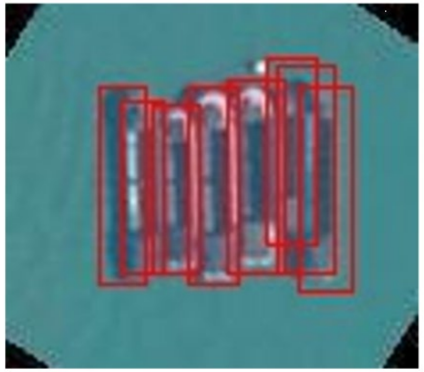

Zhang et al., 2019

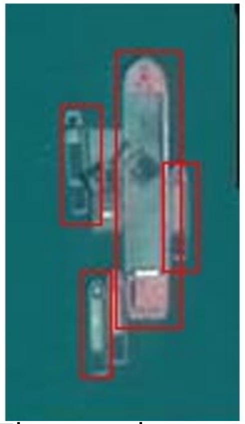

Zhang et al., 2019 v

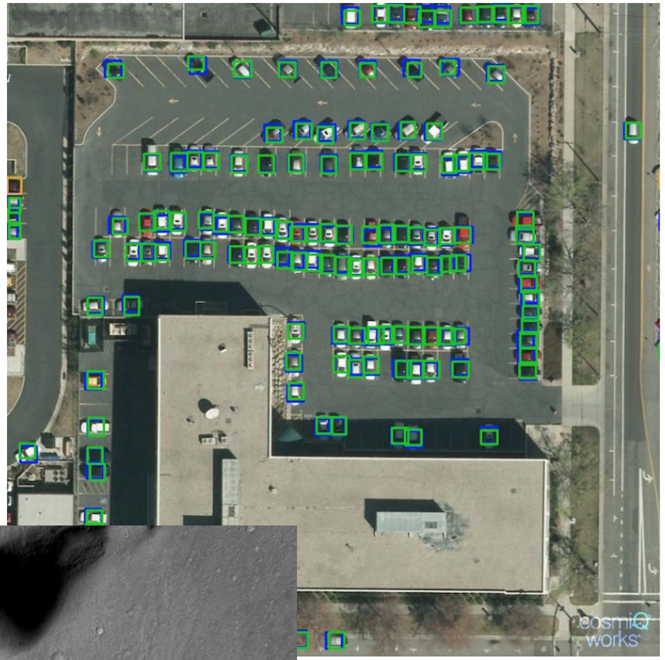

CosmiQ

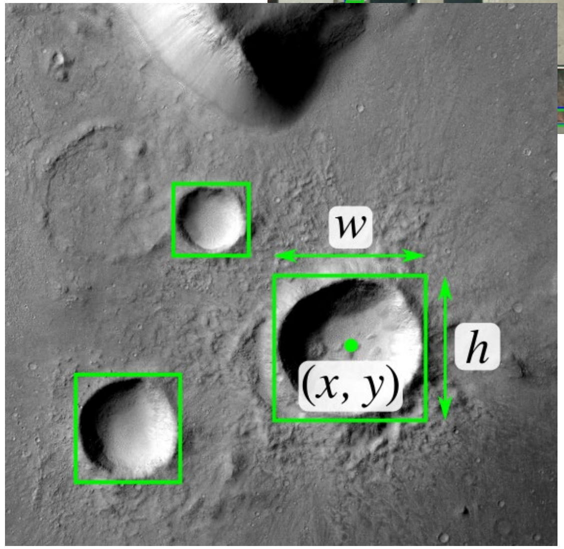

Wronkiewicz, Kerner, & Harrison 2018

Semantic segmentation

- Linking each pixel in an image to a class label
  - Supervised: class labels are known
  - Unsupervised: class labels are unknown (result is similar to clustering)
- Often involves object detection step
- Commonly deep learning architectures:
  - Fully convolutional networks (Long et al., 2015)
  - U-Net (Ronneberger et al., 2015)
  - Mask R-CNN (He et al., 2017)

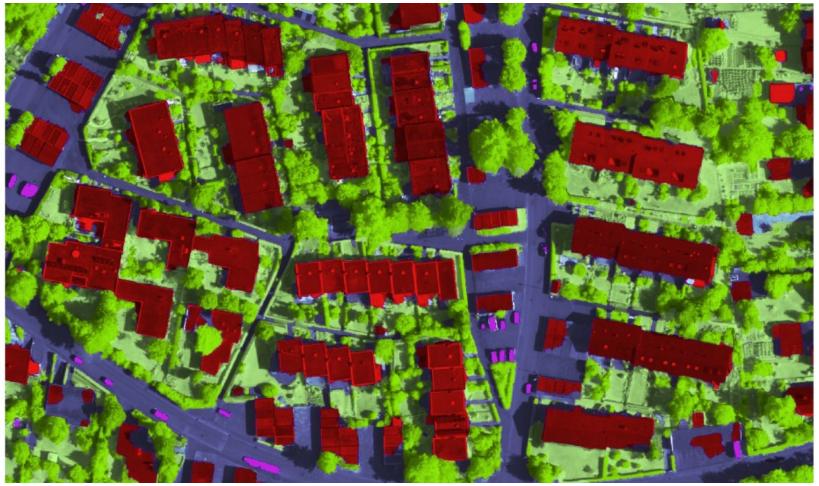

Marmanis et al., 2016

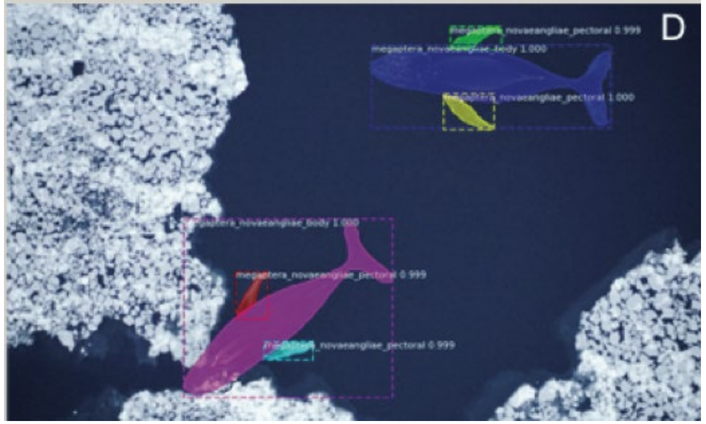

Gray et al., 2019

Pansharpening/super-resolution

- Image fusion or direct prediction of higher-resolution image from lower-resolution image
- Deep learning approaches:
  - Convolutional neural networks
  - Autoencoder neural networks
  - Generative adversarial networks (GANs)
- May be useful for object detection or combining data sources, but concern about physical interpretation of predicted data
  - “Deep fakes”: https://thispersondoesnotexist.com/

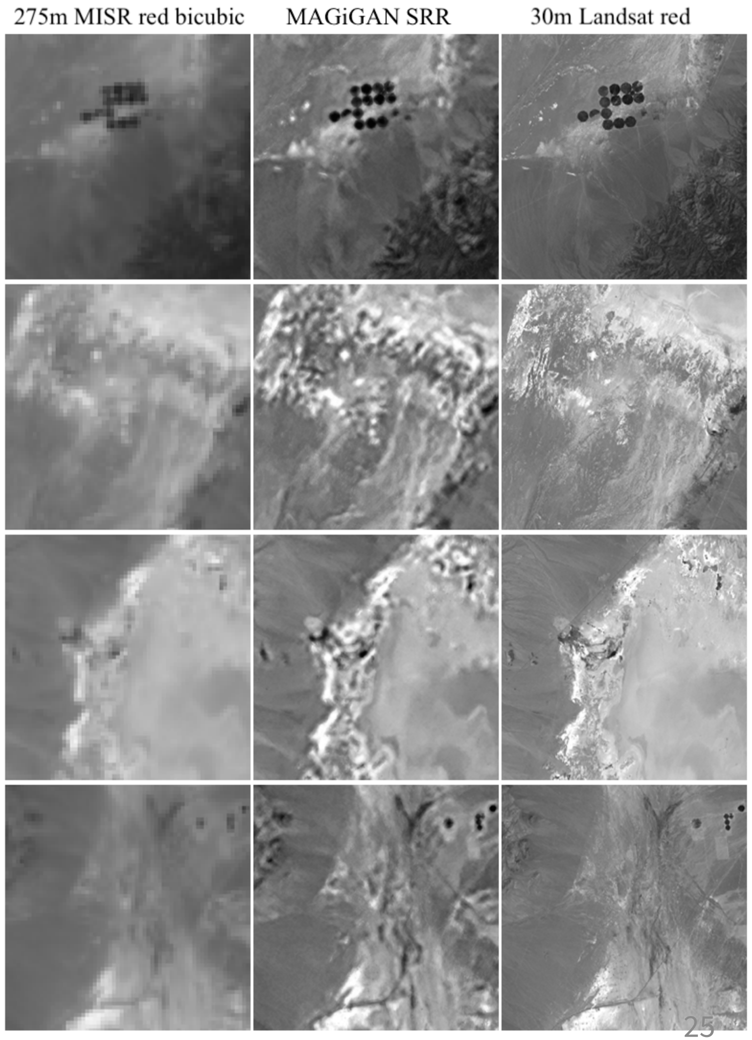

Tao et al., 2019

Image Registration

- Aligning two or more images captured at different times/viewpoints/sensors
  - Goal to find affine transformation between pair of images
- Pre-processing step for many approaches, e.g., change detection

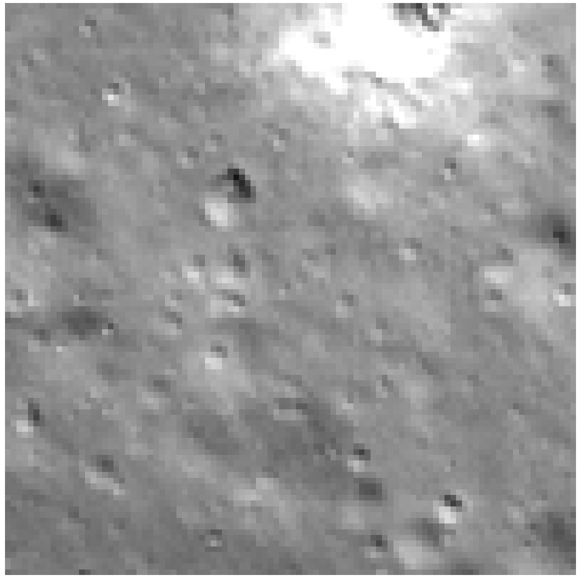

Mis-registered LROC images
Kerner et al., 2019
Credit: LROC/ASU

Image Registration

- Aligning two or more images captured at different times/viewpoints/sensors
  - Goal to find affine transformation between pair of images
- Pre-processing step for many approaches, e.g., change detection
- Many approaches based on Siamese networks
- Others estimate image-to-image mapping function using regression methods including random forest or standard neural networks

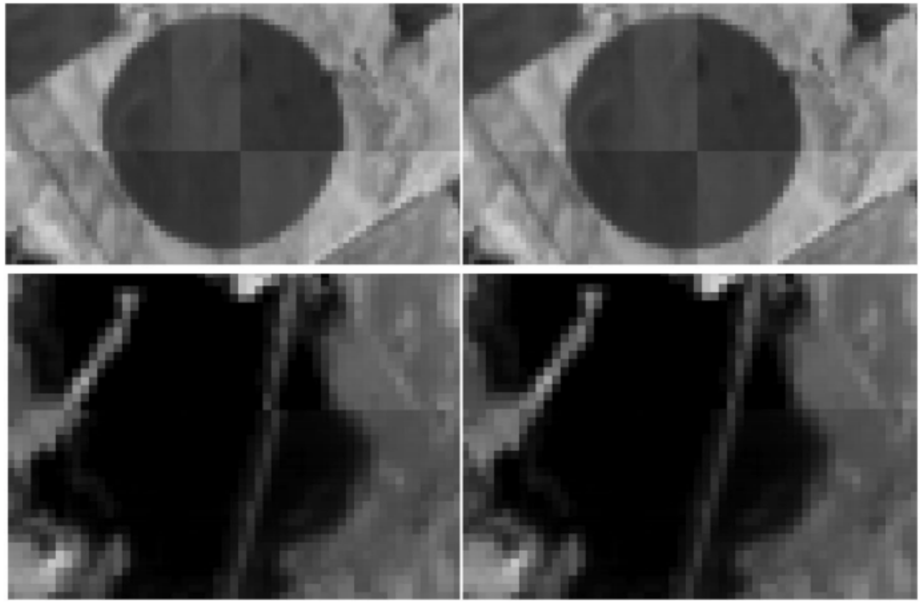

Zoomed-in “chessboard” of Landsat-8 and Sentinel-2A images before and after co-registration using Random Forest regression mapping.
(Skakun et al., 2017)

Limitations/Directions for Future Work

- Land cover classification most common, but little discussion of generalization
    - Most studies focused on producing one map for one year in one area
- Global inference/operational analysis is still a major challenge
- Limited ground data, especially for developing countries without national programs
- Interpretability/explainability and reproducibility
    - Uncommon for studies to go "extra step" to understand results and learned representations or provide executable code

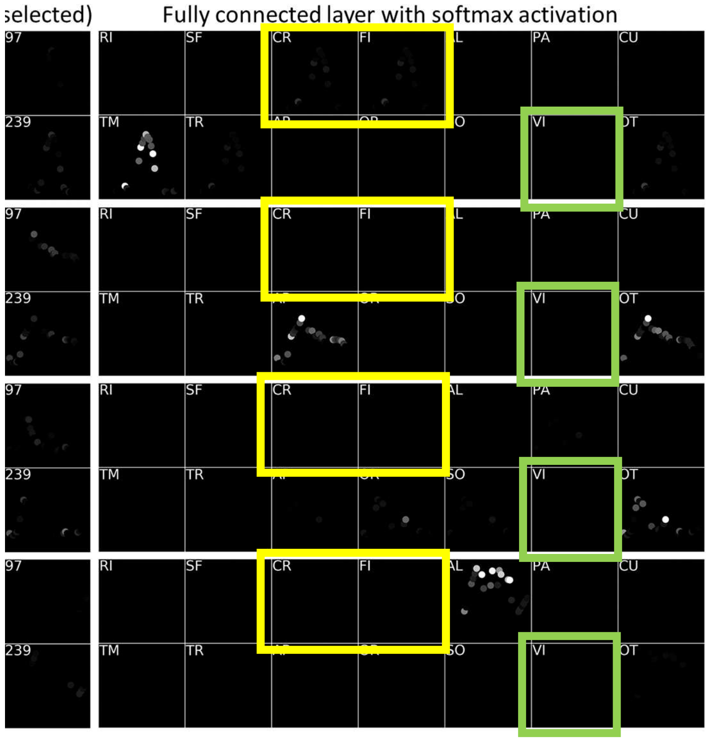

Visualization of feature map activations for different EVI time series input examples. (Zhong et al., 2019)

Local peaks

Decreasing slope

Questions?

HARVEST

EARTH DATA

FOR INFORMED AGRICULTURAL DECISIONS

Hannah Kerner

Assistant Research Professor

University of Maryland, College Park

hkerner@umd.edu