Tell Me Why: Interpretable Machine Learning for Space Exploration and Beyond

Kiri L. Wagstaff[1], Jake Lee[1], and Nina Lanza[2]

[1]Jet Propulsion Laboratory, California Institute of Technology

[2]Los Alamos National Laboratory

ASU SESE Colloquium, October 31, 2018

This work was performed in part at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA. CL#18-6078

Discovery in Large Data Sets

Scientific discoveries often come from outliers

Flickr user Klau


Brandon Dilbeck

Ted Hood, State Library of NSW

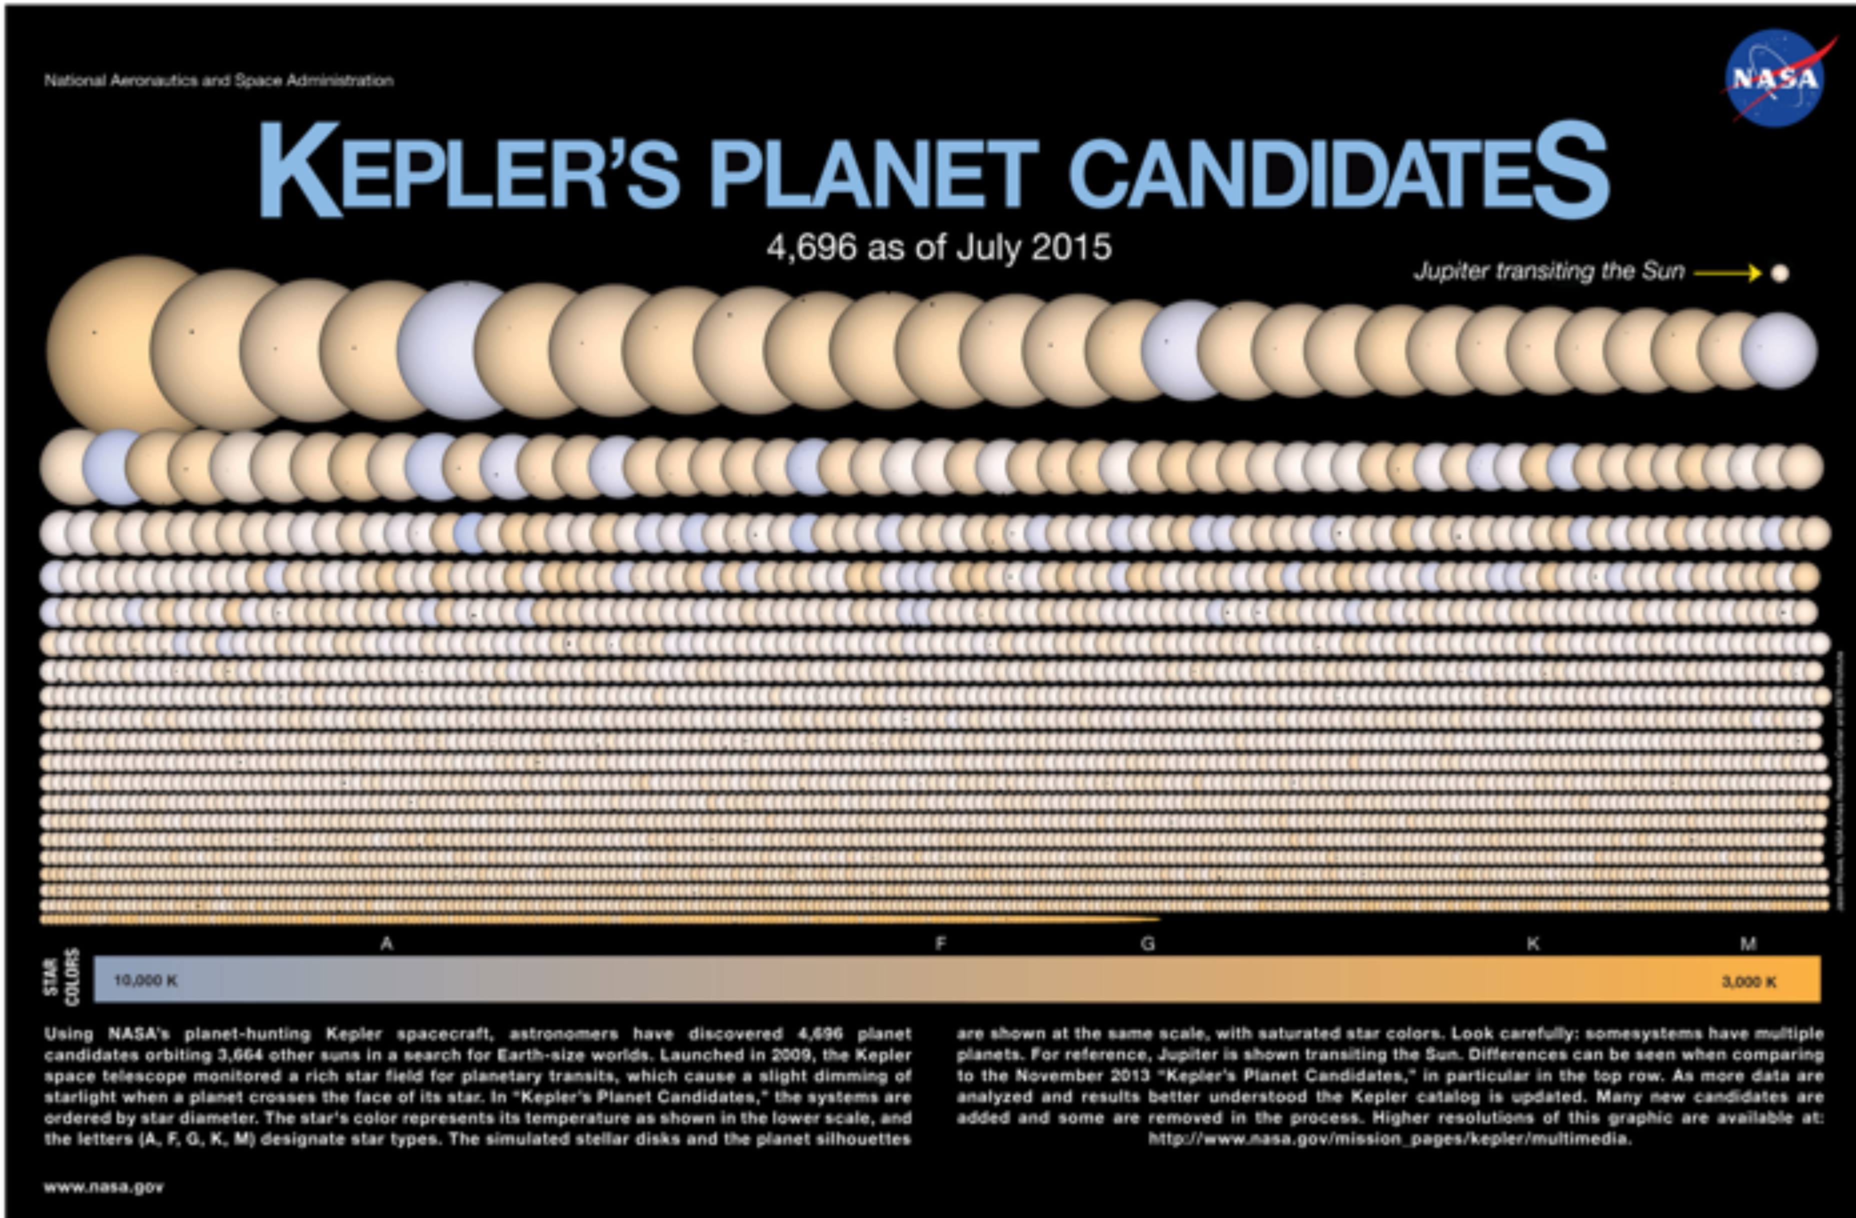

National Aeronautics and Space Administration
NASA
KEPLER'S PLANET CANDIDATES
4,696 as of July 2015
Jupiter transiting the Sun
STAR COLORS
A
F
G
K
M
10,000 K
3,000 K
Using NASA's planet-hunting Kepler spacecraft, astronomers have discovered 4,696 planet candidates orbiting 3,664 other suns in a search for Earth-size worlds. Launched in 2009, the Kepler space telescope monitored a rich star field for planetary transits, which cause a slight dimming of starlight when a planet crosses the face of its star. In "Kepler's Planet Candidates," the systems are ordered by star diameter. The star's color represents its temperature as shown in the lower scale, and the letters (A, F, G, K, M) designate star types. The simulated stellar disks and the planet silhouettes are shown at the same scale, with saturated star colors. Look carefully: somesystems have multiple planets. For reference, Jupiter is shown transiting the Sun. Differences can be seen when comparing to the November 2013 "Kepler's Planet Candidates," in particular in the top row. As more data are analyzed and results better understood the Kepler catalog is updated. Many new candidates are added and some are removed in the process. Higher resolutions of this graphic are available at: http://www.nasa.gov/mission_pages/kepler/multimedia.
www.nasa.gov

Machine Learning

Algorithms that learn a concept model from examples

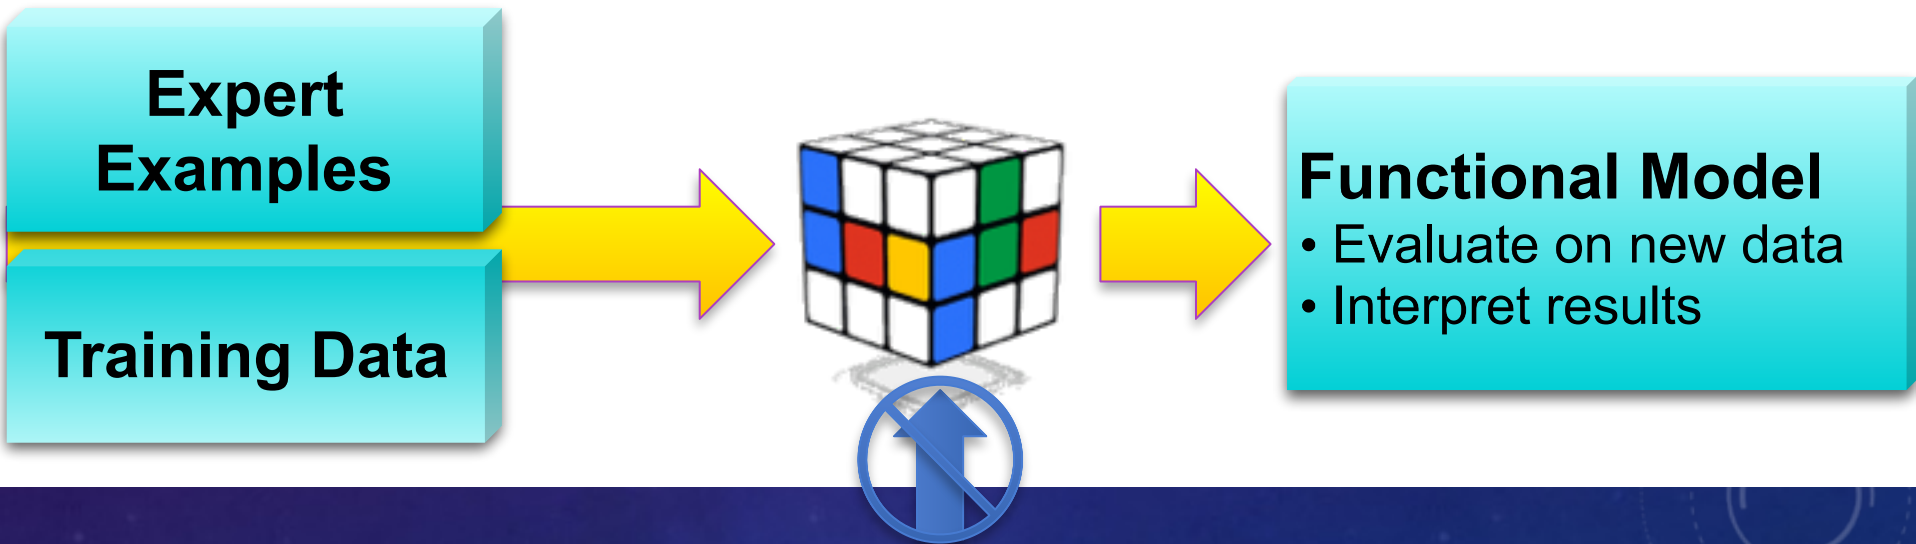

Strengths:

No need to specify rules

No need to explain “how”

Machine Learning for Novelty Detection

DEMUD: Discovery via Eigenbasis Modeling of Uninteresting Data

- Prioritizes interesting observations within large data sets
- Minimizes redundancy in selections (models what you already know)
- Provides explanations for why items are chosen

ChemCam on Mars

- Laser + spectrometer = remote composition
- 5,860 wavelengths
- 4,238 targets in 707 sols
  - Excludes calibration targets and poor-quality data

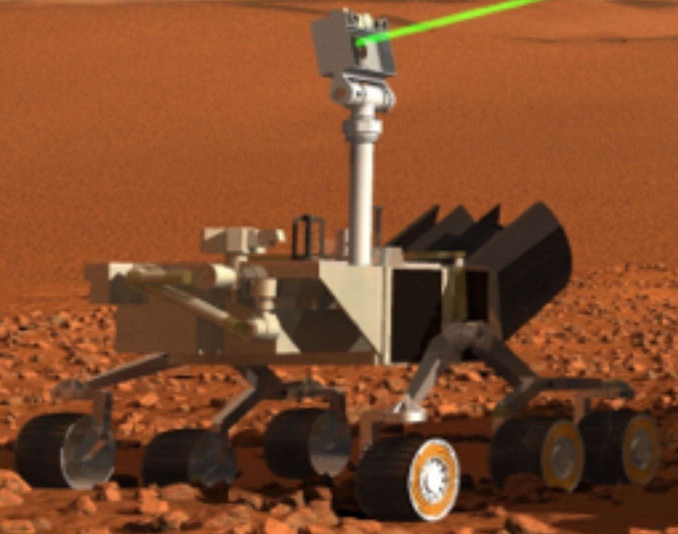

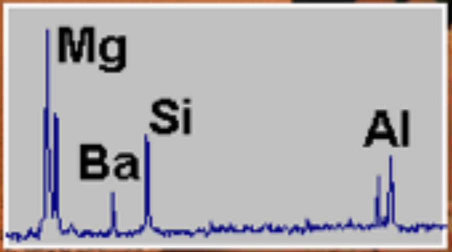

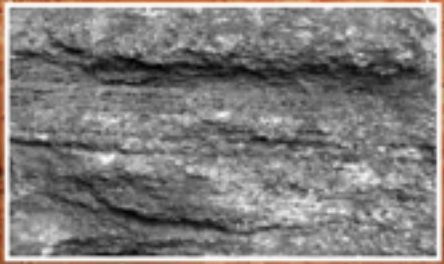

ChemCam Observations

Epworth2

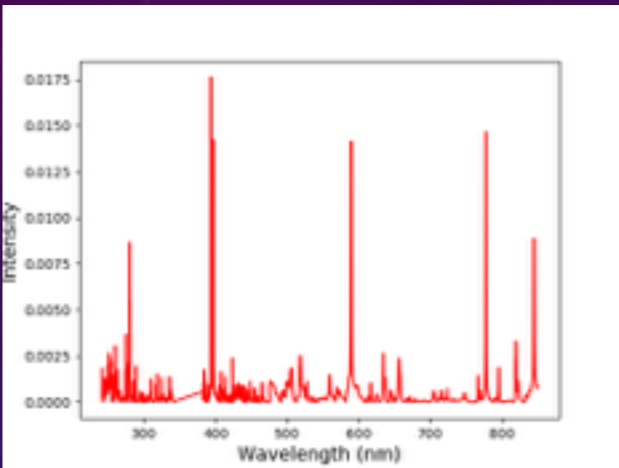

Rocknest6b

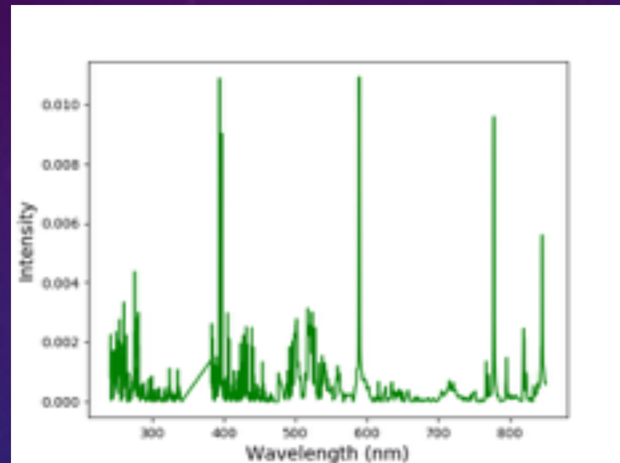

Kenyon

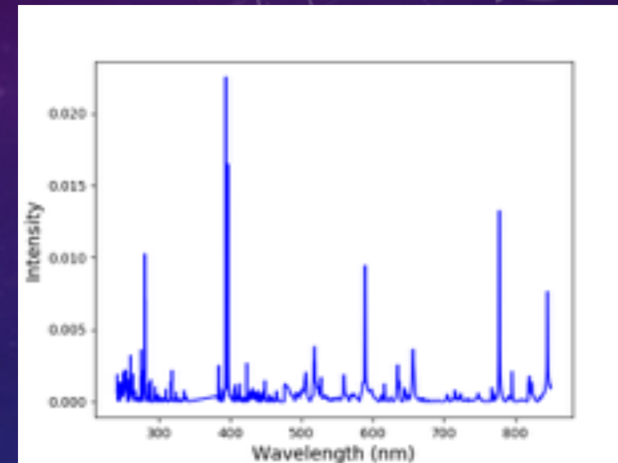

Preble

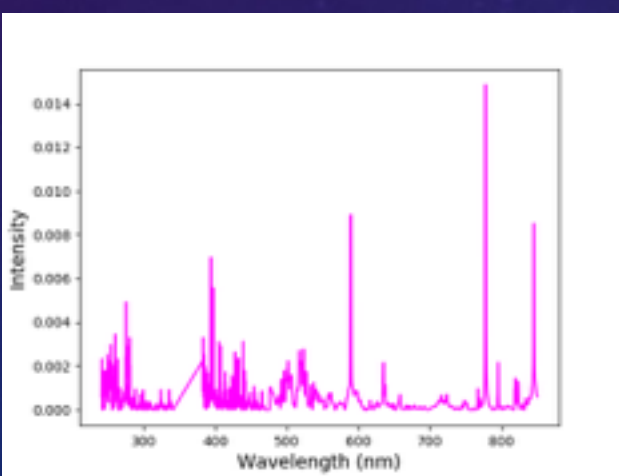

Pearson

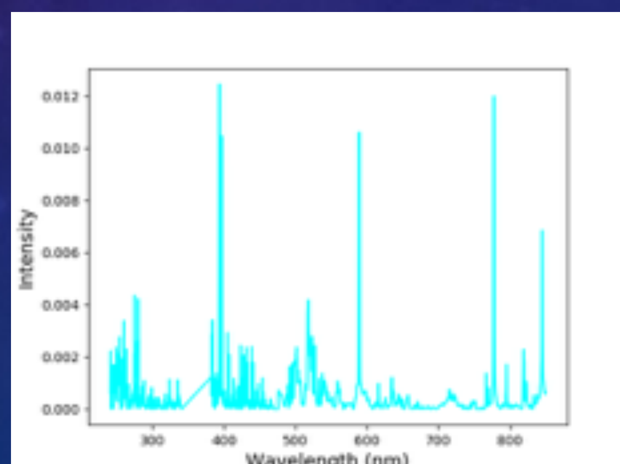

Rocknest7

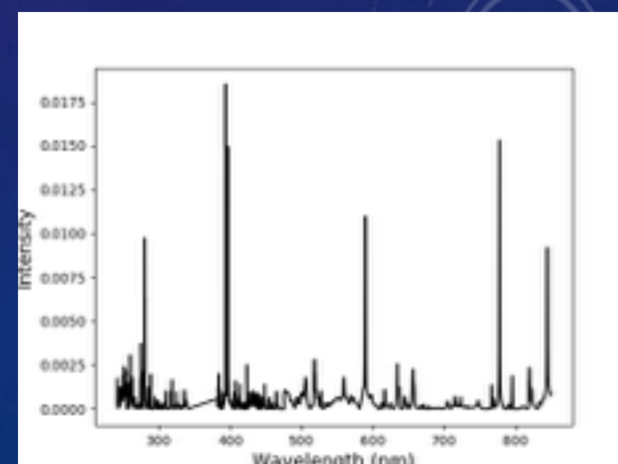

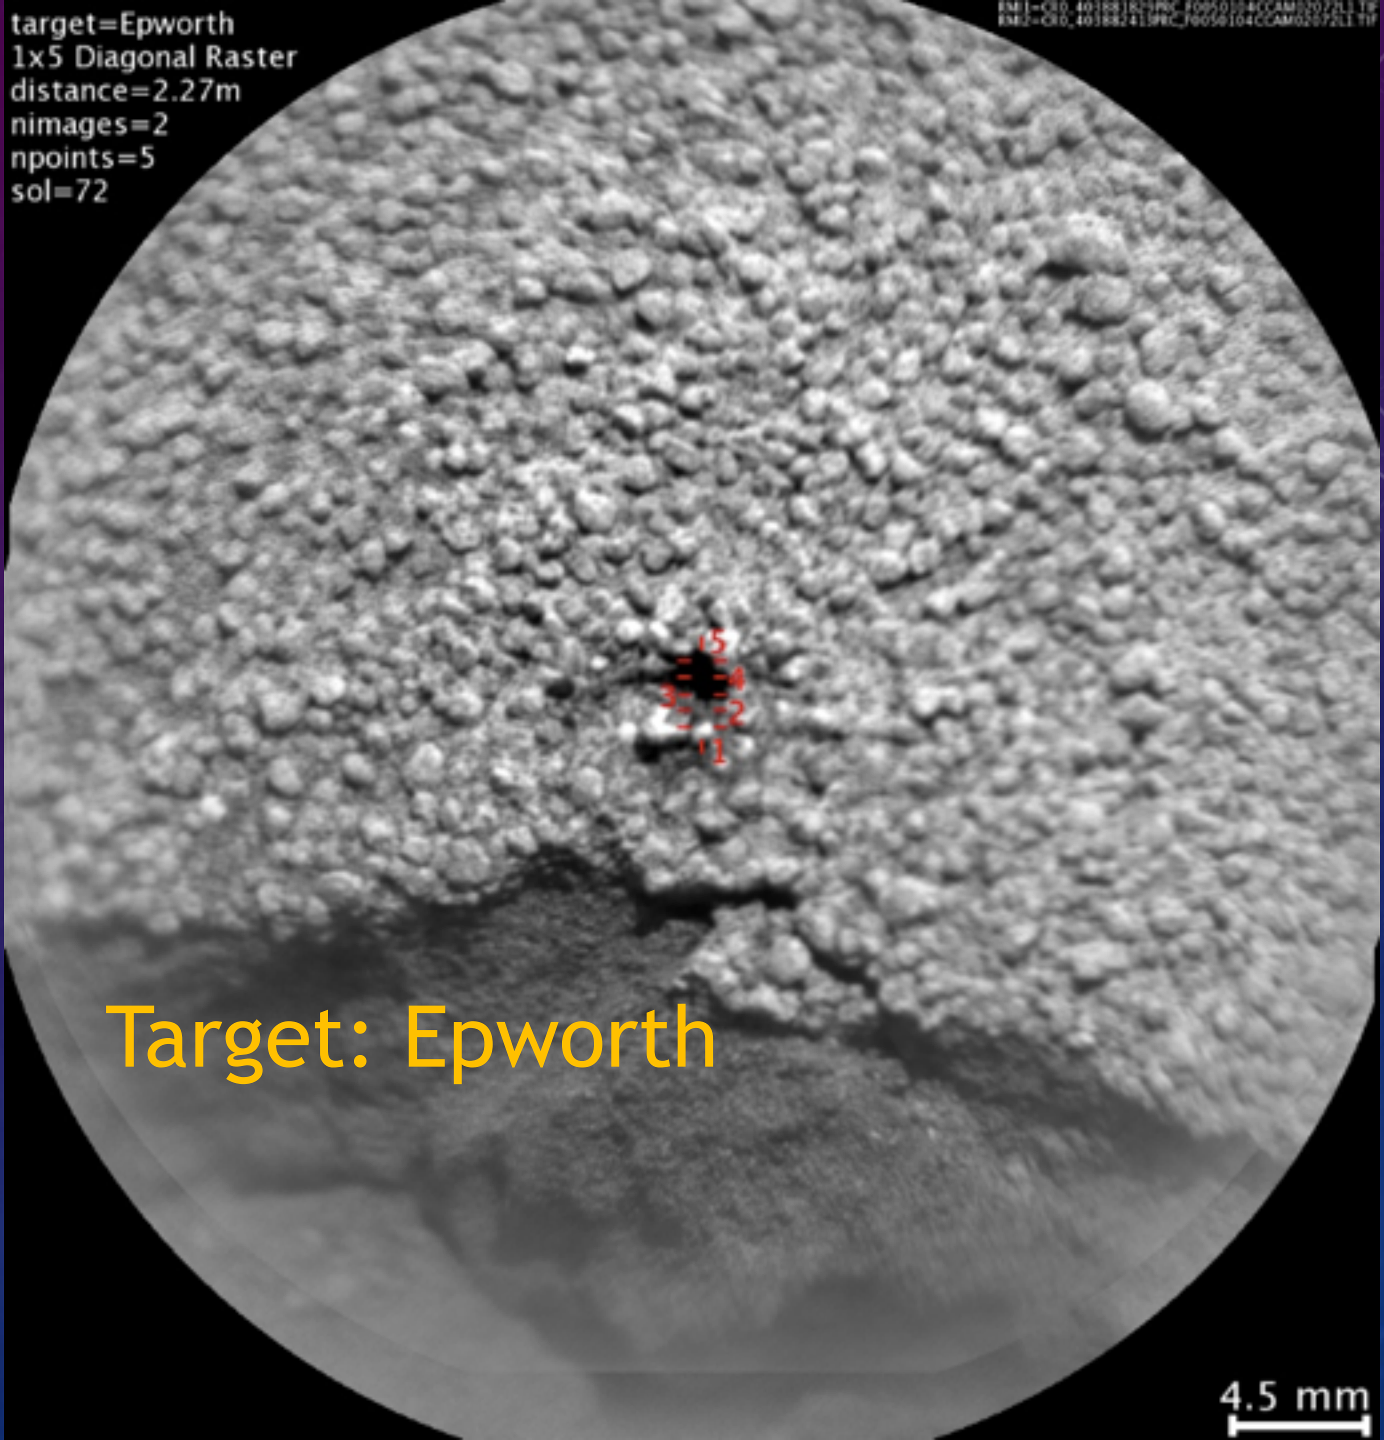

Target: Epworth

DEMUD's view of Epworth

[Wagstaff et al., 2014]

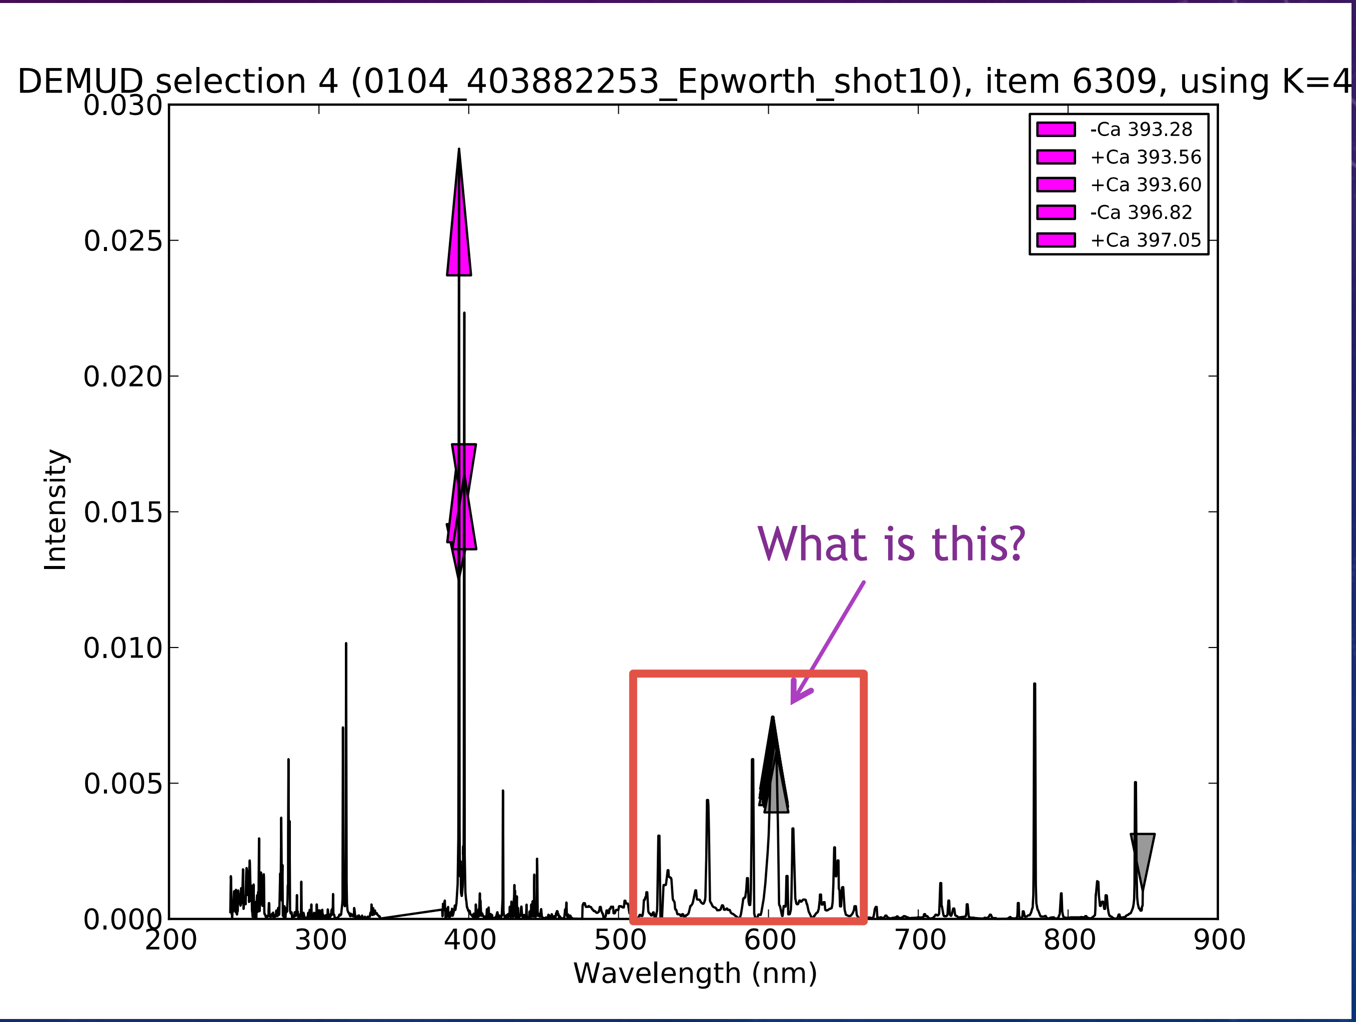

“First fluorine detection on Mars with ChemCam on-board MSL”

[Forni et al., 2014]

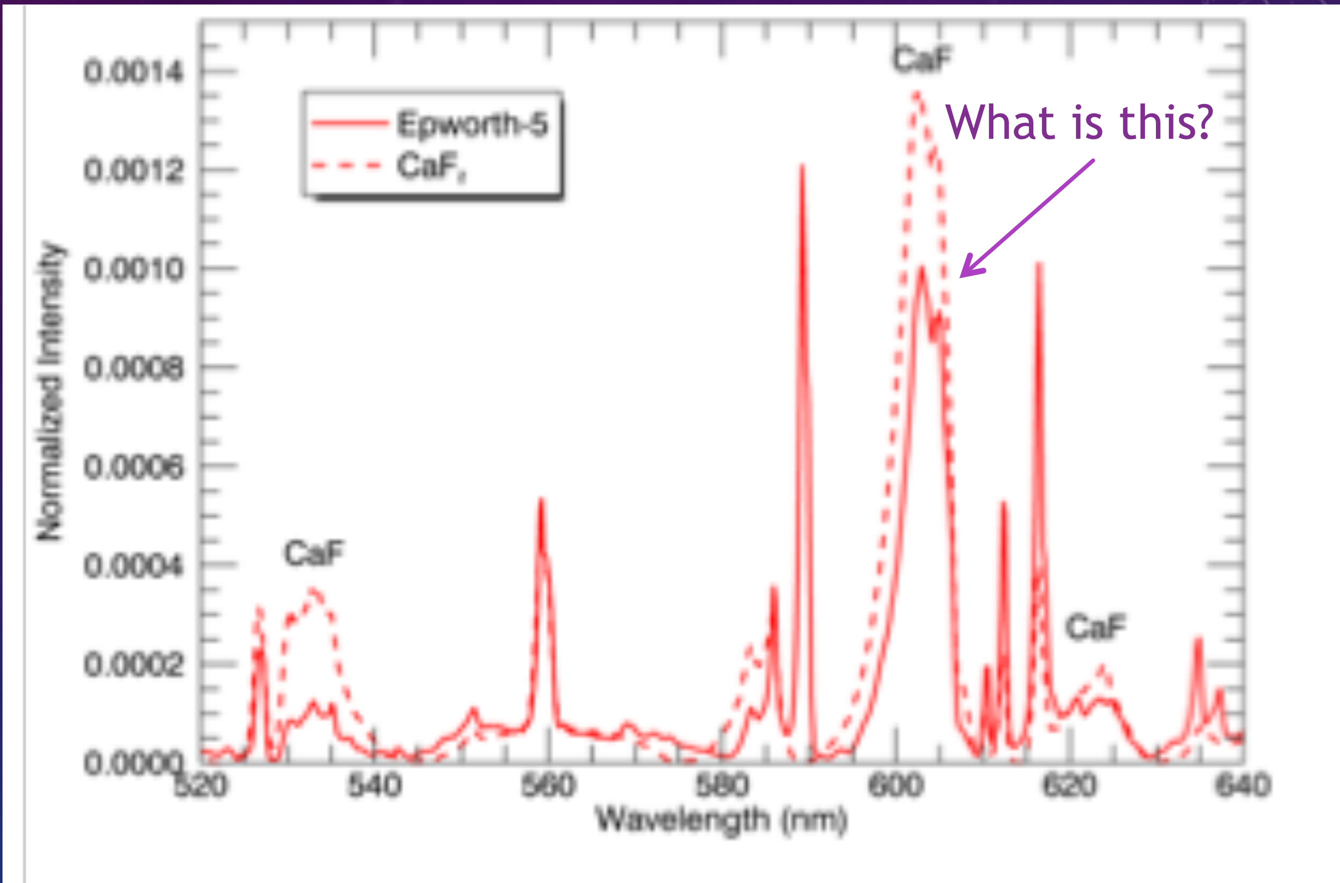

Figure 1. Comparison of Epworth-5 spectrum with a pure CaF_2 laboratory spectrum obtained under Martian conditions

Traditional Anomaly Detection

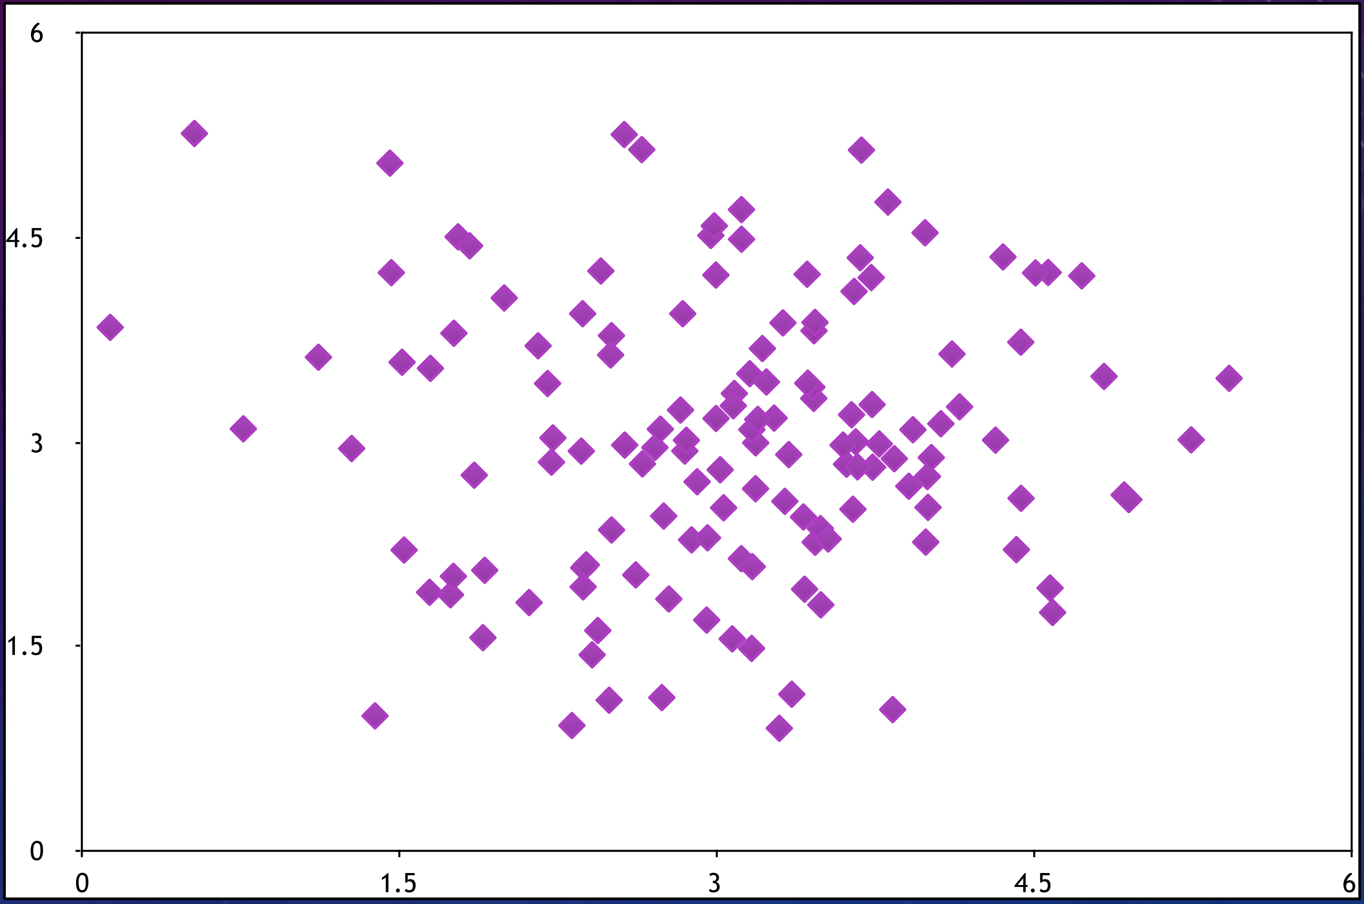

Traditional Anomaly Detection

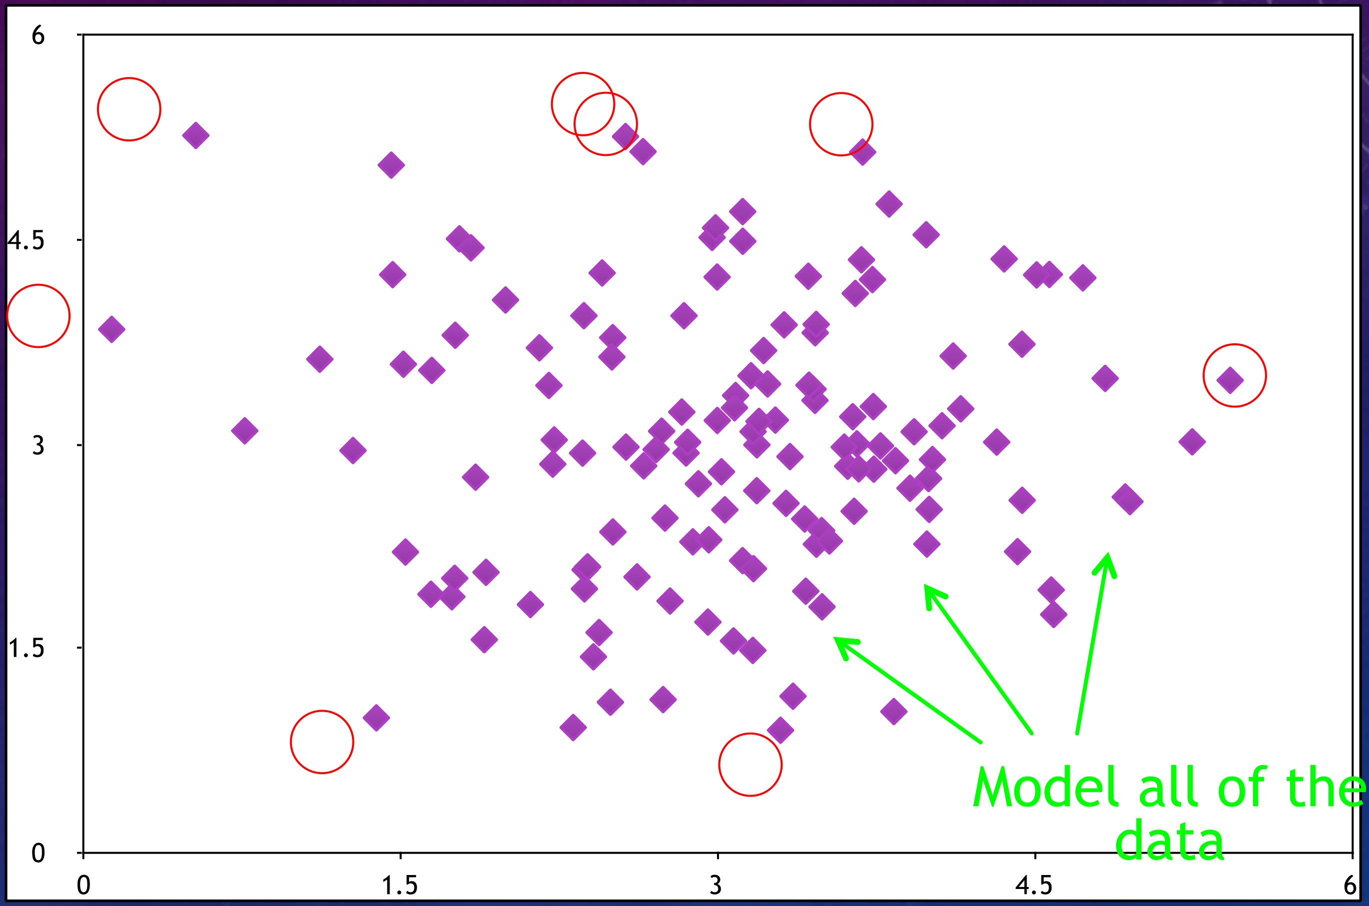

DEMUD Novelty Detection

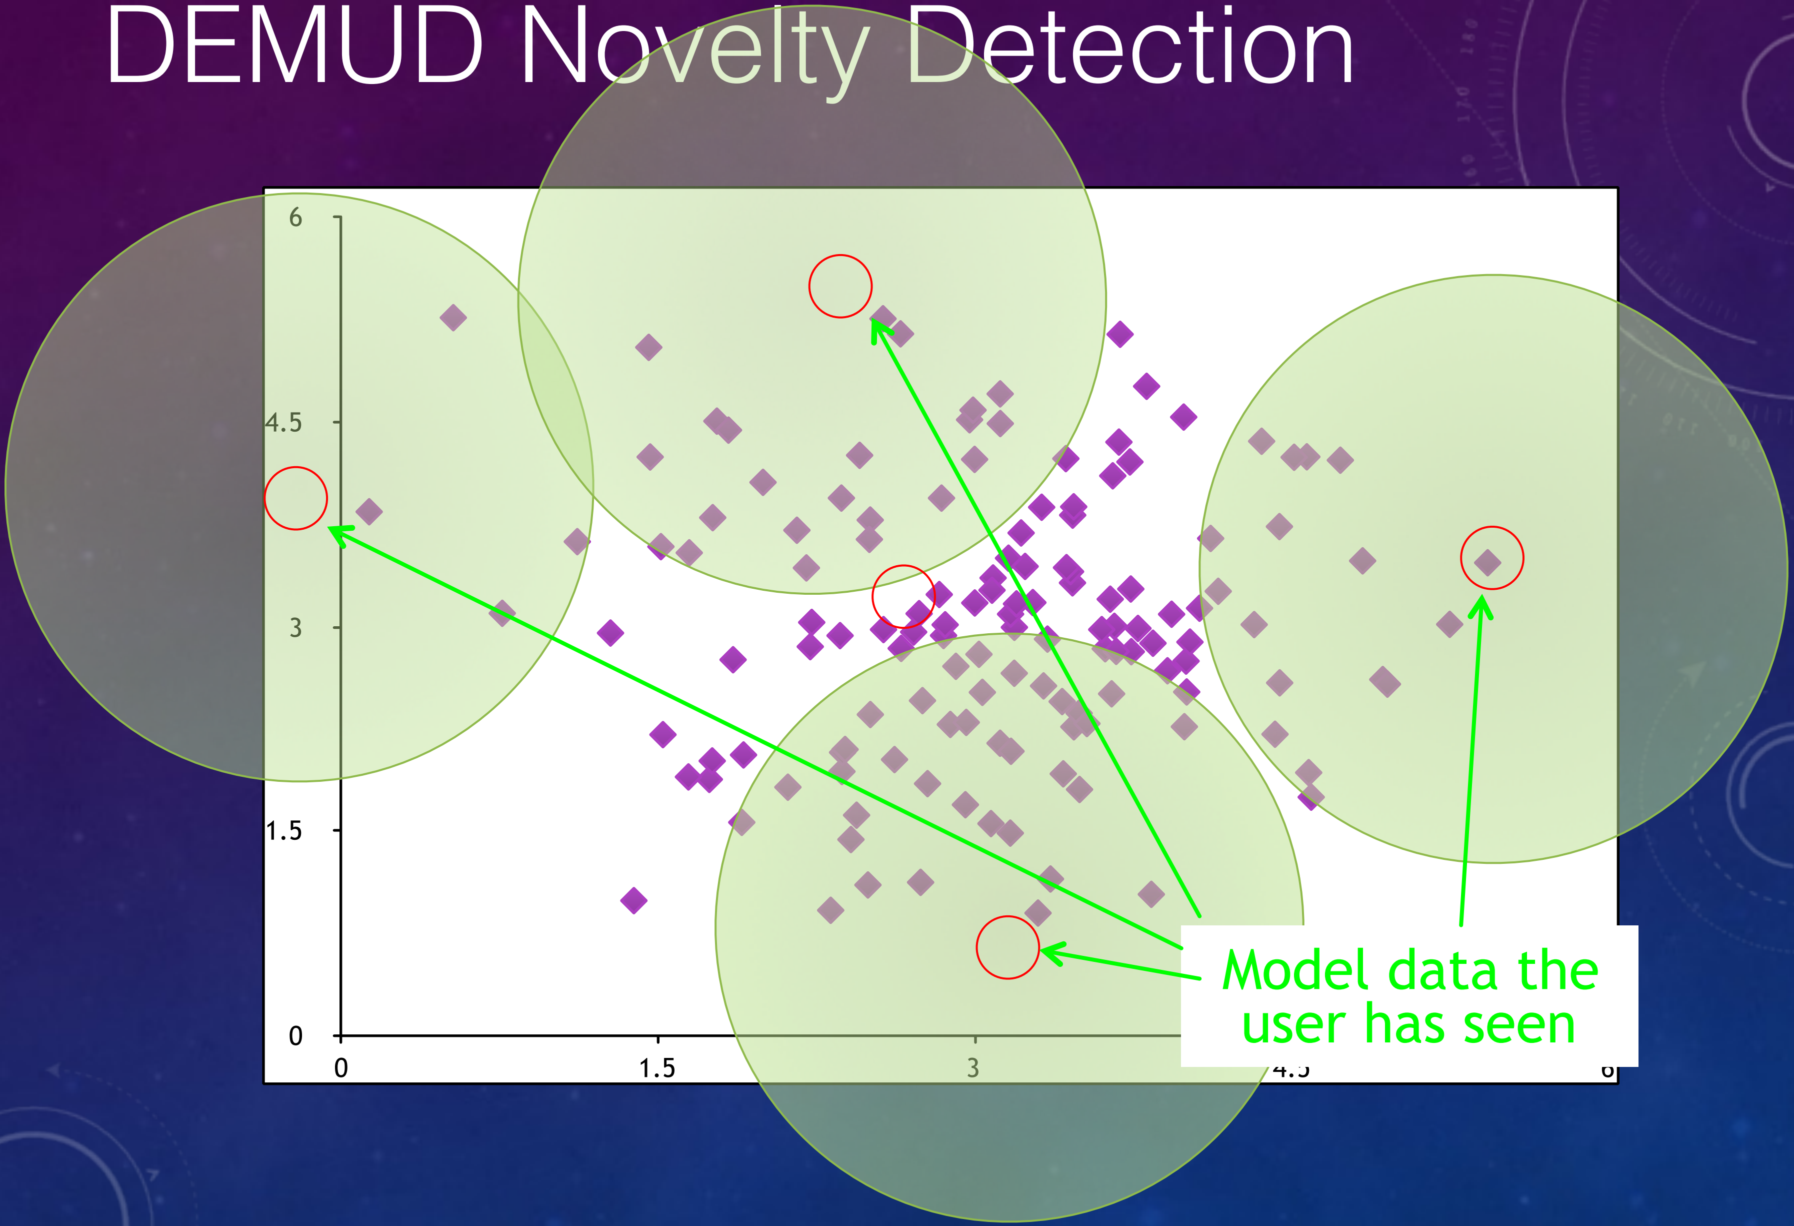

Novelty Detection: DEMUD

- Incremental discovery using SVD model of selections
  - Build a model of selected items X *not the entire data set D*
  - Select new items that are difficult to represent with the model
  - Explanation = information that is new
  - Update model with each selection

Model learns along with you

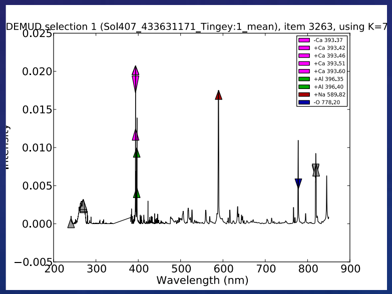

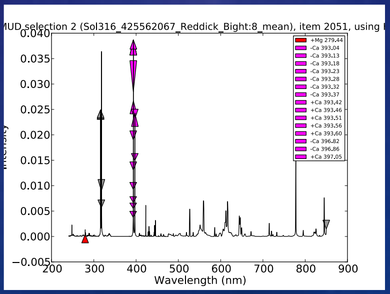

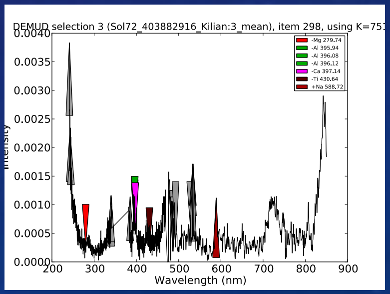

…

Discovery in Large Image Data Sets

Surveillance

Human faces
Credit: Pixabay user Geralt

Planetary Science

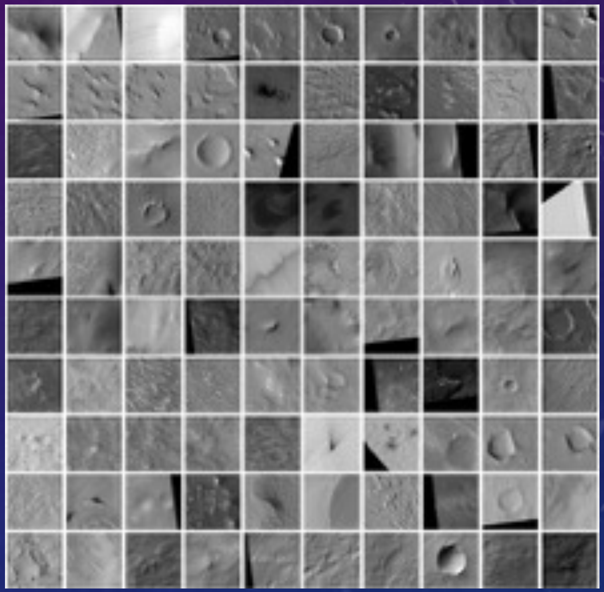

HiRISE – Mars surface features from orbit
Credit: NASA/JPL-Caltech/Univ. of Arizona

- PDS Imaging Node: > 1 PB of image data

DEMUD for Images

- Representation
  - Raw pixels

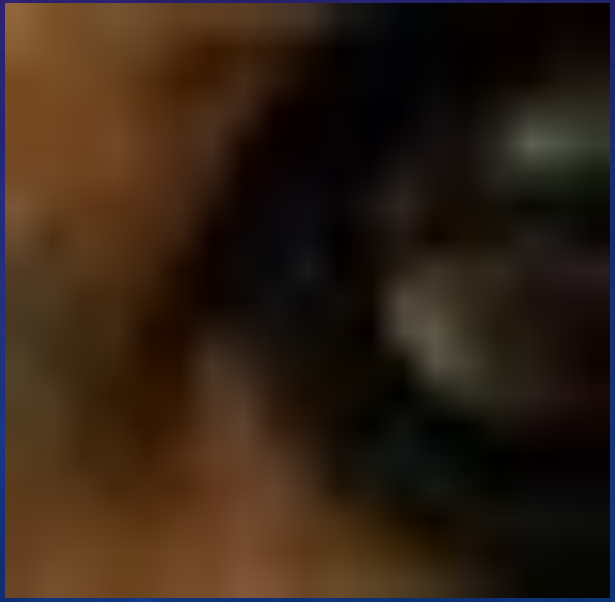

DEMUD for Images

- Representation
  - Raw pixels
  - SIFT [Lowe, 2004]

DEMUD for Images

- Representation
  - Raw pixels
  - SIFT [Lowe, 2004]
  - Neural network features [Razavian et al., 2014]

DEMUD + Neural Network Representations

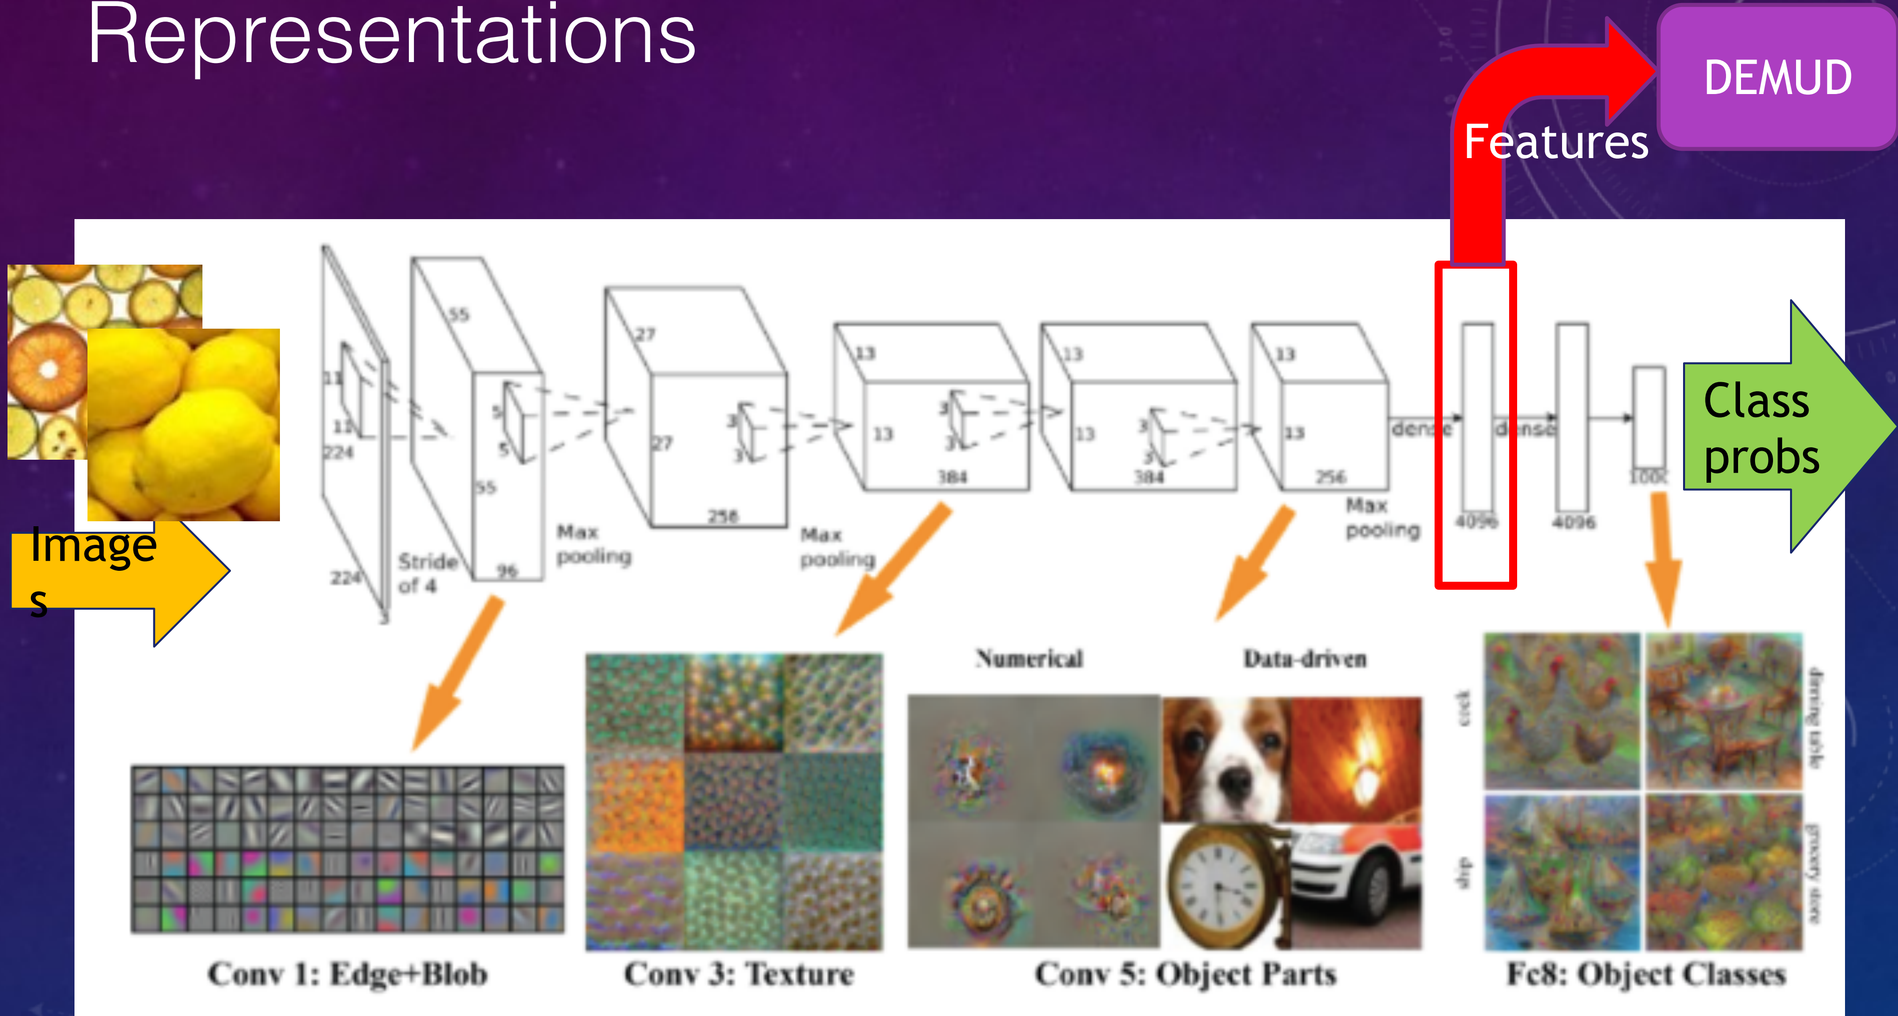

[Wei et al.]

Explanations with Neural Network Features

Selection

Features

DEMUD

4096

Explanations with Neural Network Features

Features

4096

DEMUD

Selection

4096

CNN Feature Inversion:

Predict original image with second NN

(Dos 016)

Explanations with Neural Network Features

Features

4096

DEMUD

Selection

4096

CNN Feature Inversion:

Predict original image with second NN

(Dos 016)

Explanation

4096

?

Explanations are Context-Dependent

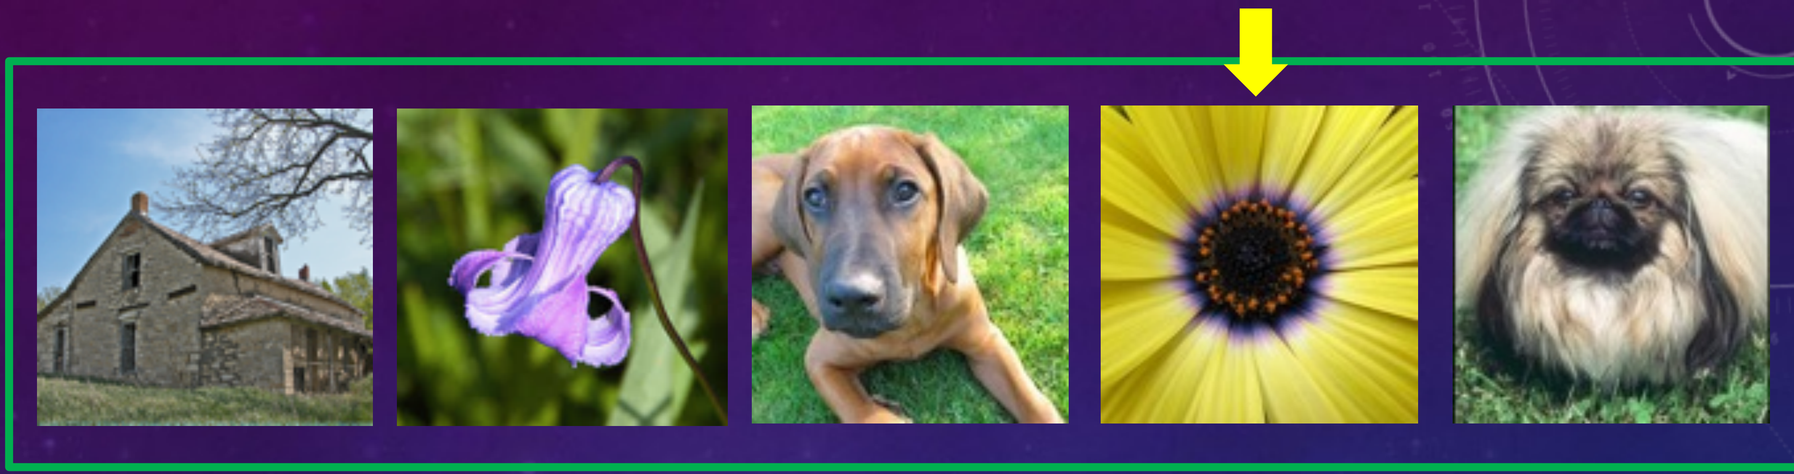

Explanations with Neural Network Features

Features

4096

DEMUD

Selection

4096

CNN Feature Inversion:

Predict original image with second NN

(Dos 016)

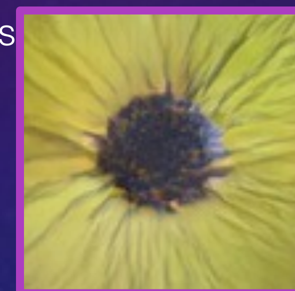

Explanation

4096

?

4096

Already

Explanations with Neural Network Features

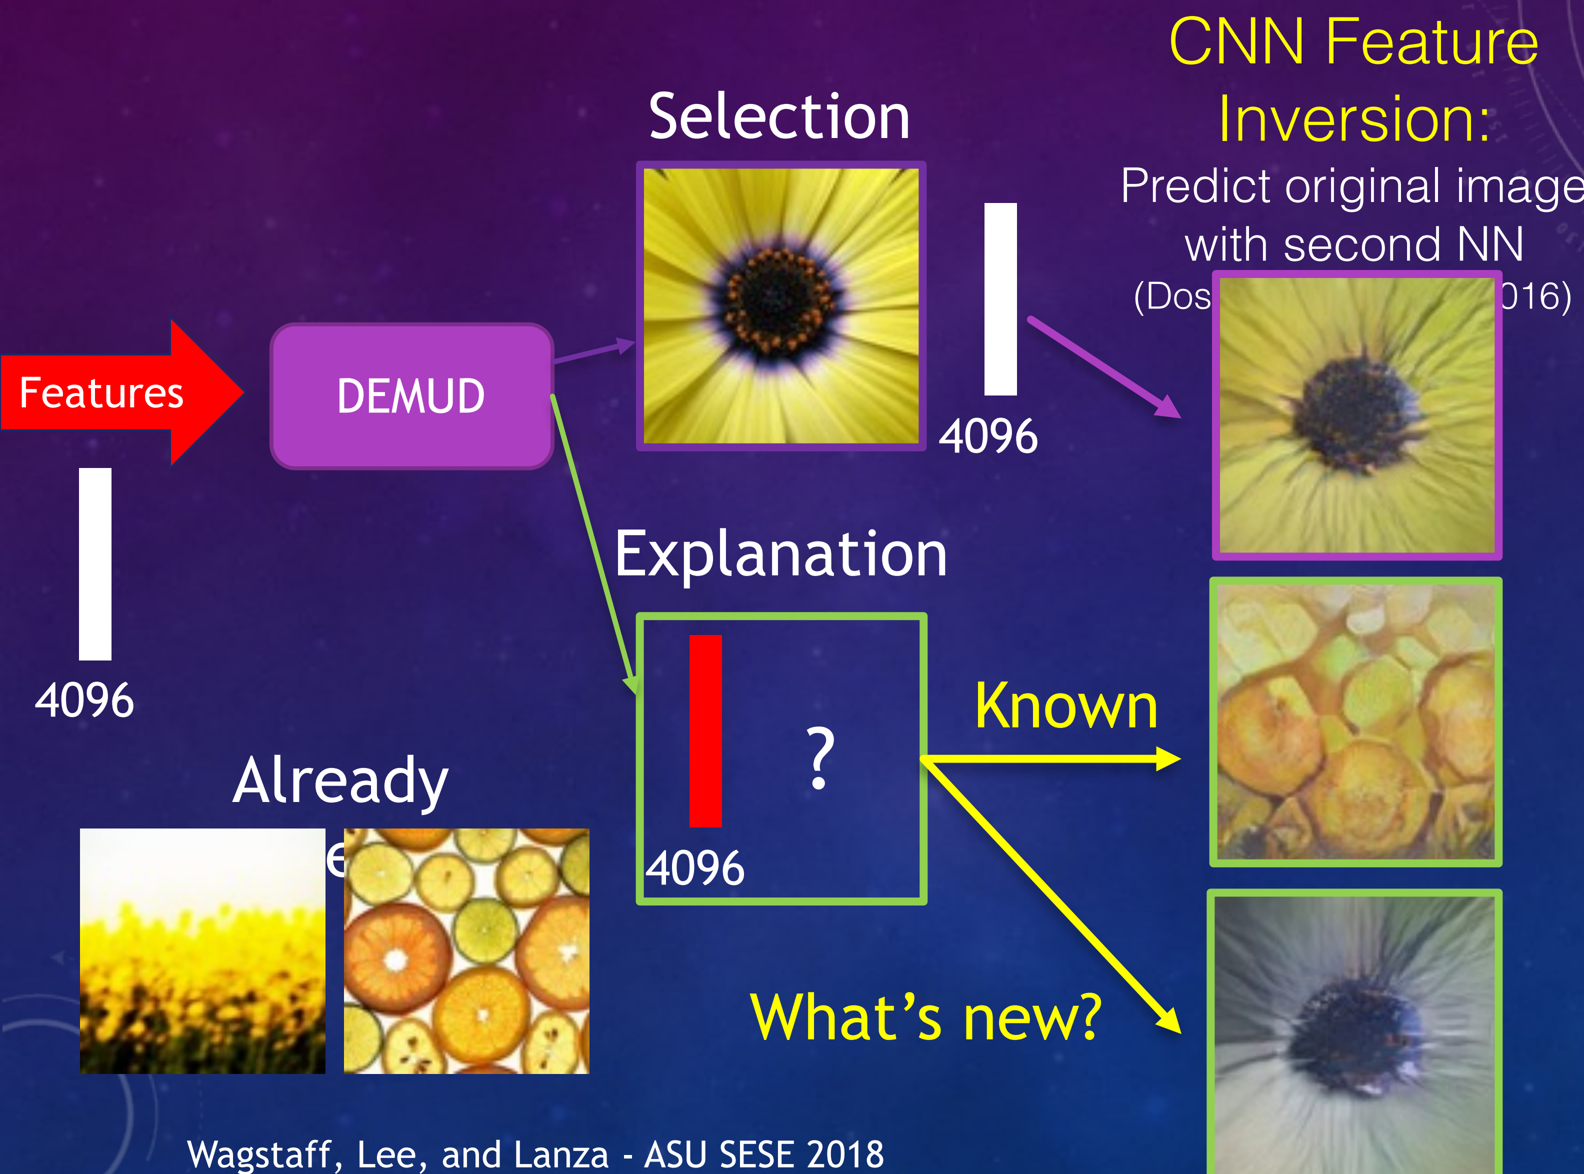

Interpretable Image Discovery

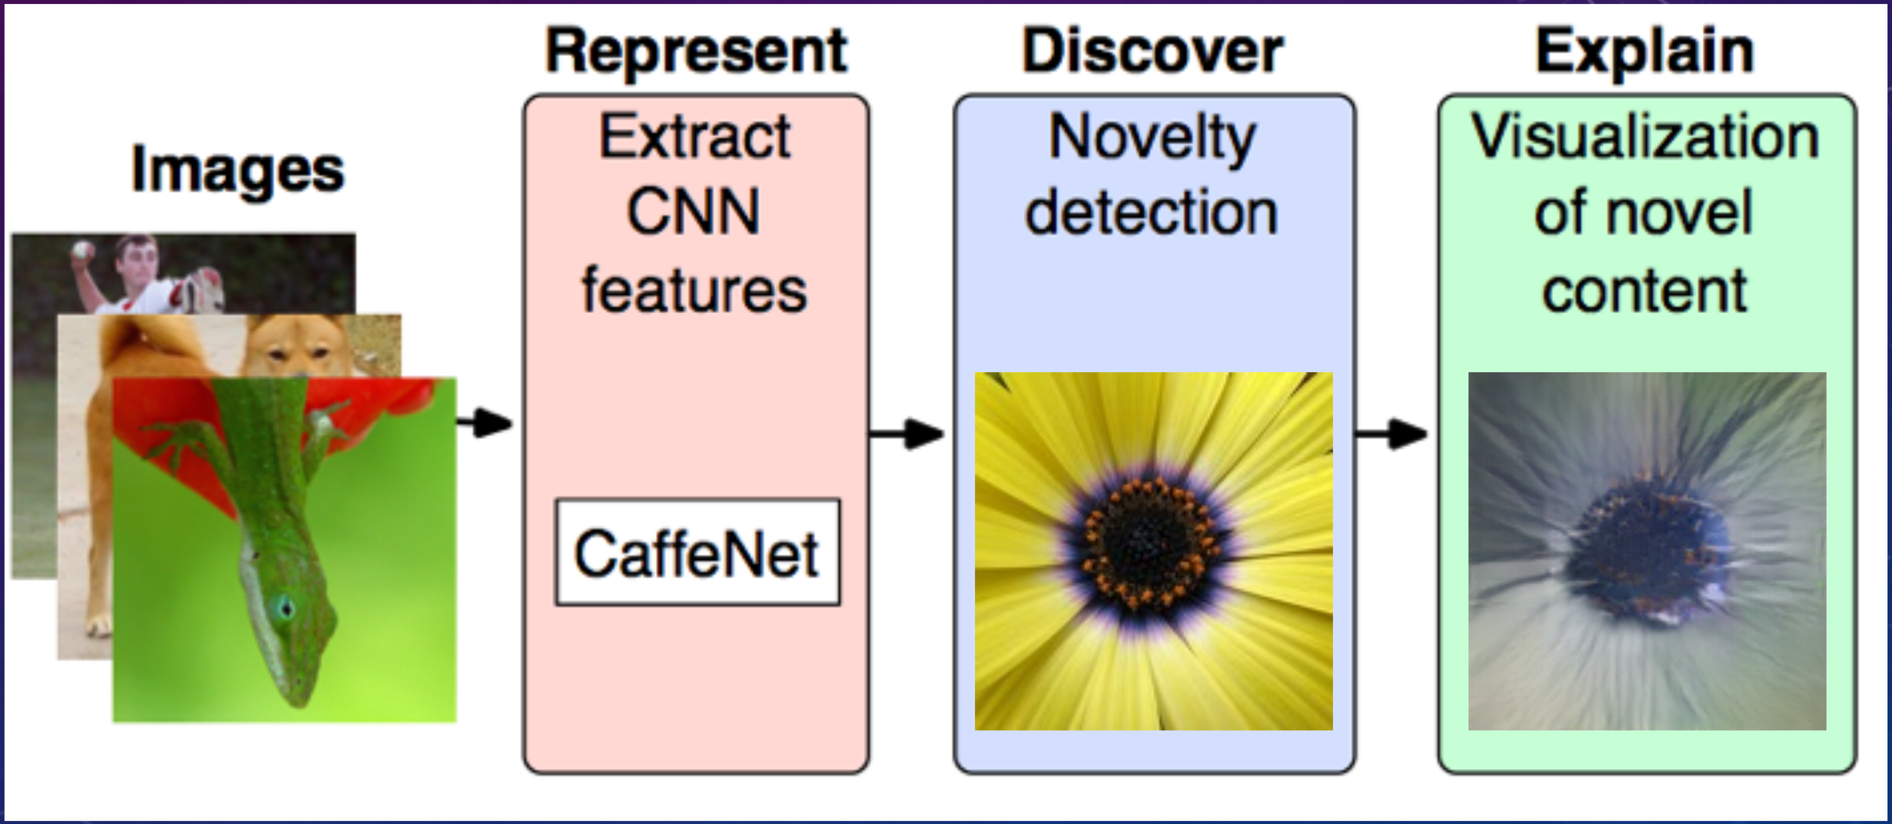

Experiments – MSL Rover images

- 6737 images: Mastcam, Navcam, MAHLI
- 25 classes: rover parts, ground, h sun
- Uneven distribution
- CNN was trained Earth images; can it help here?

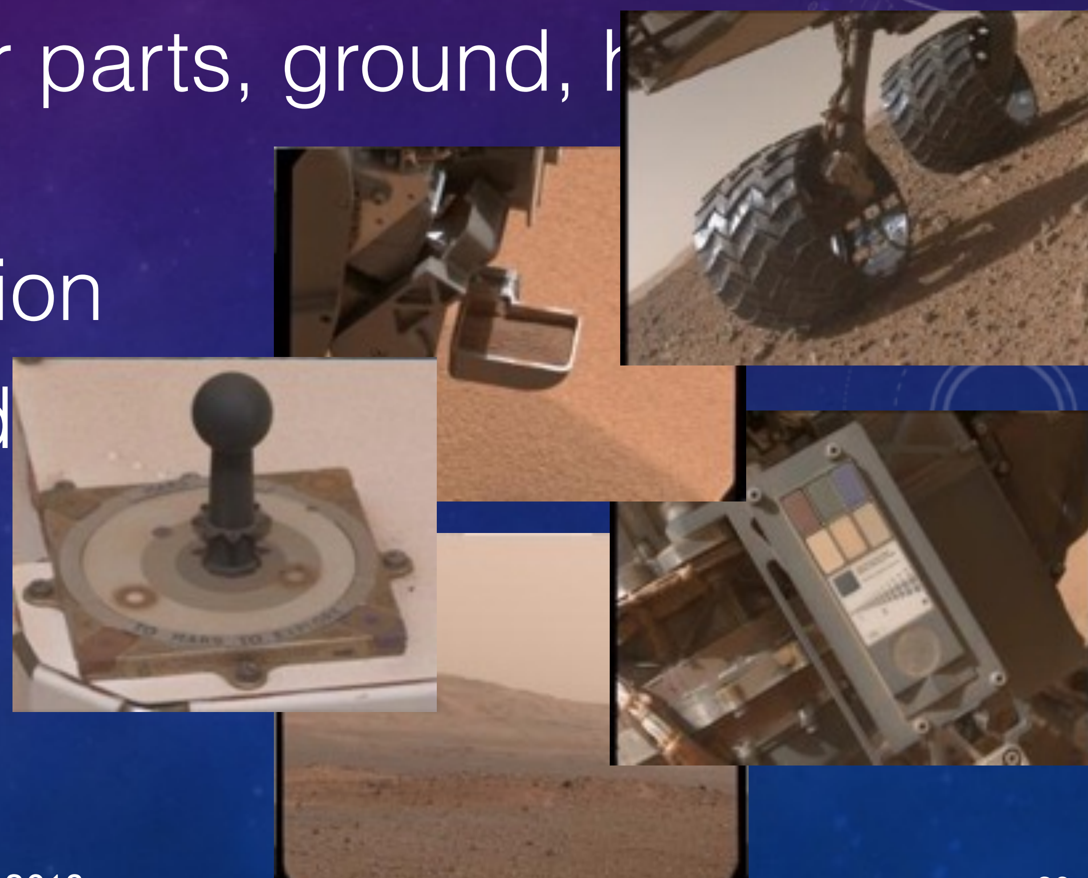

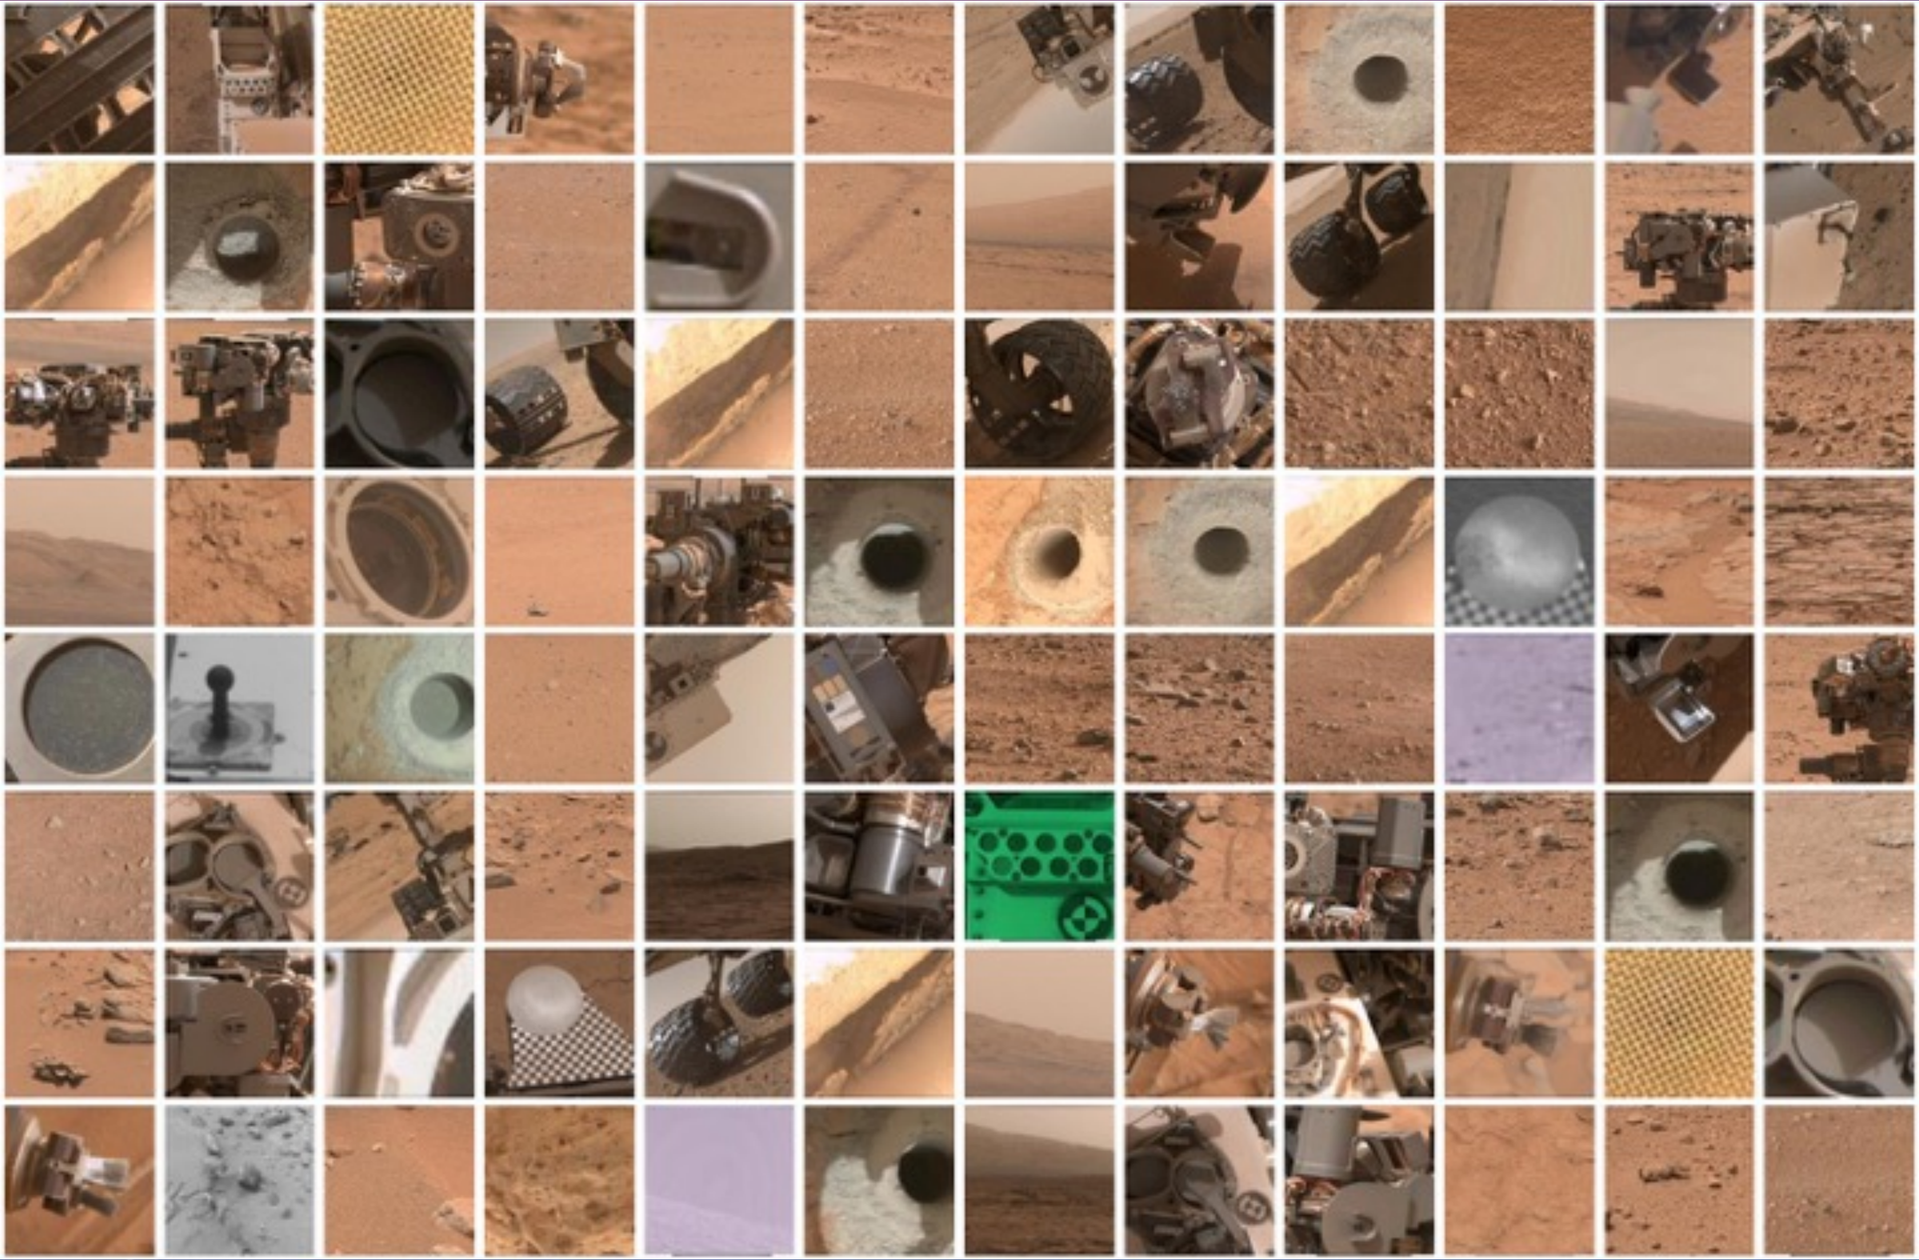

Experiments – MSL Rover images

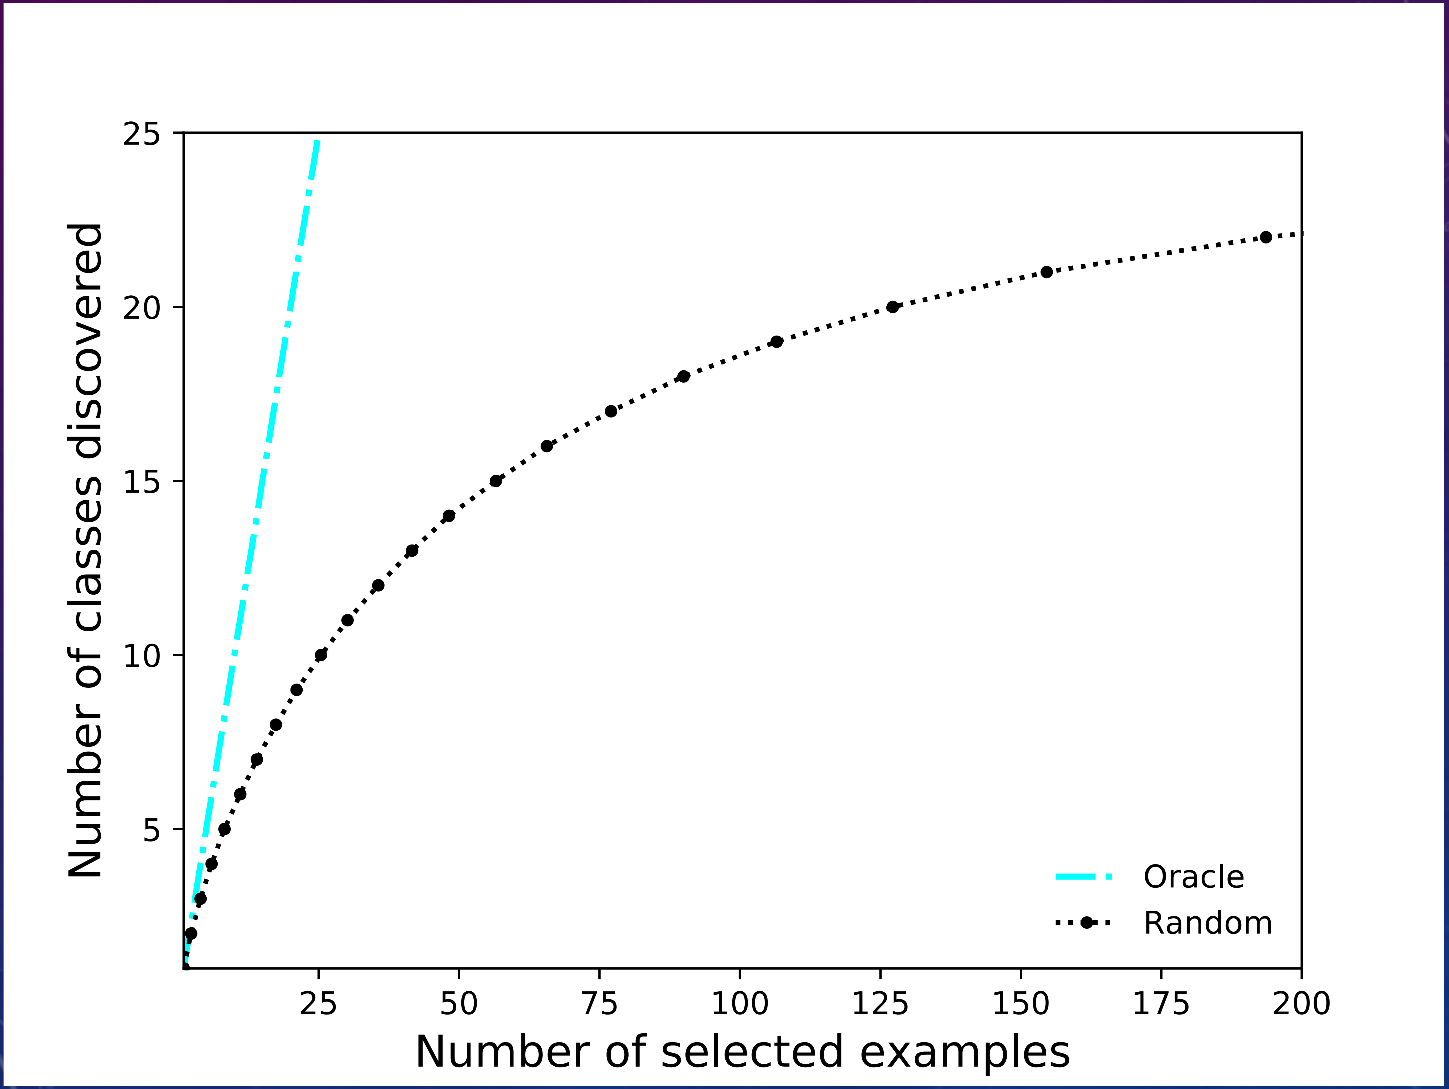

Experiments – MSL Rover images

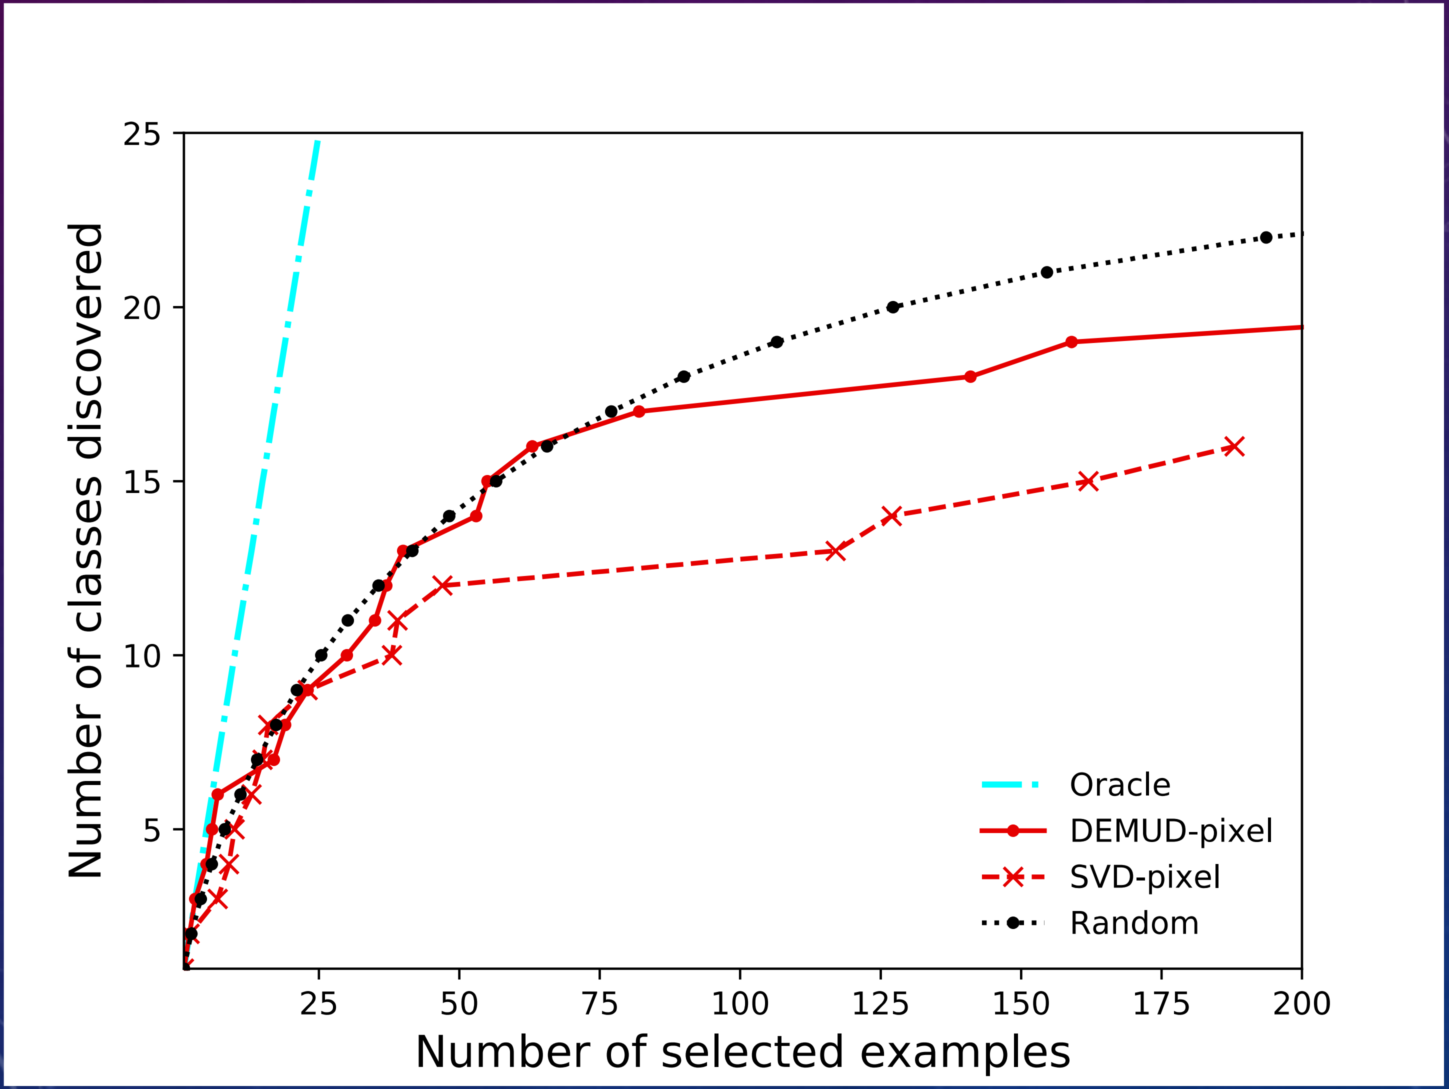

Experiments – MSL Rover images

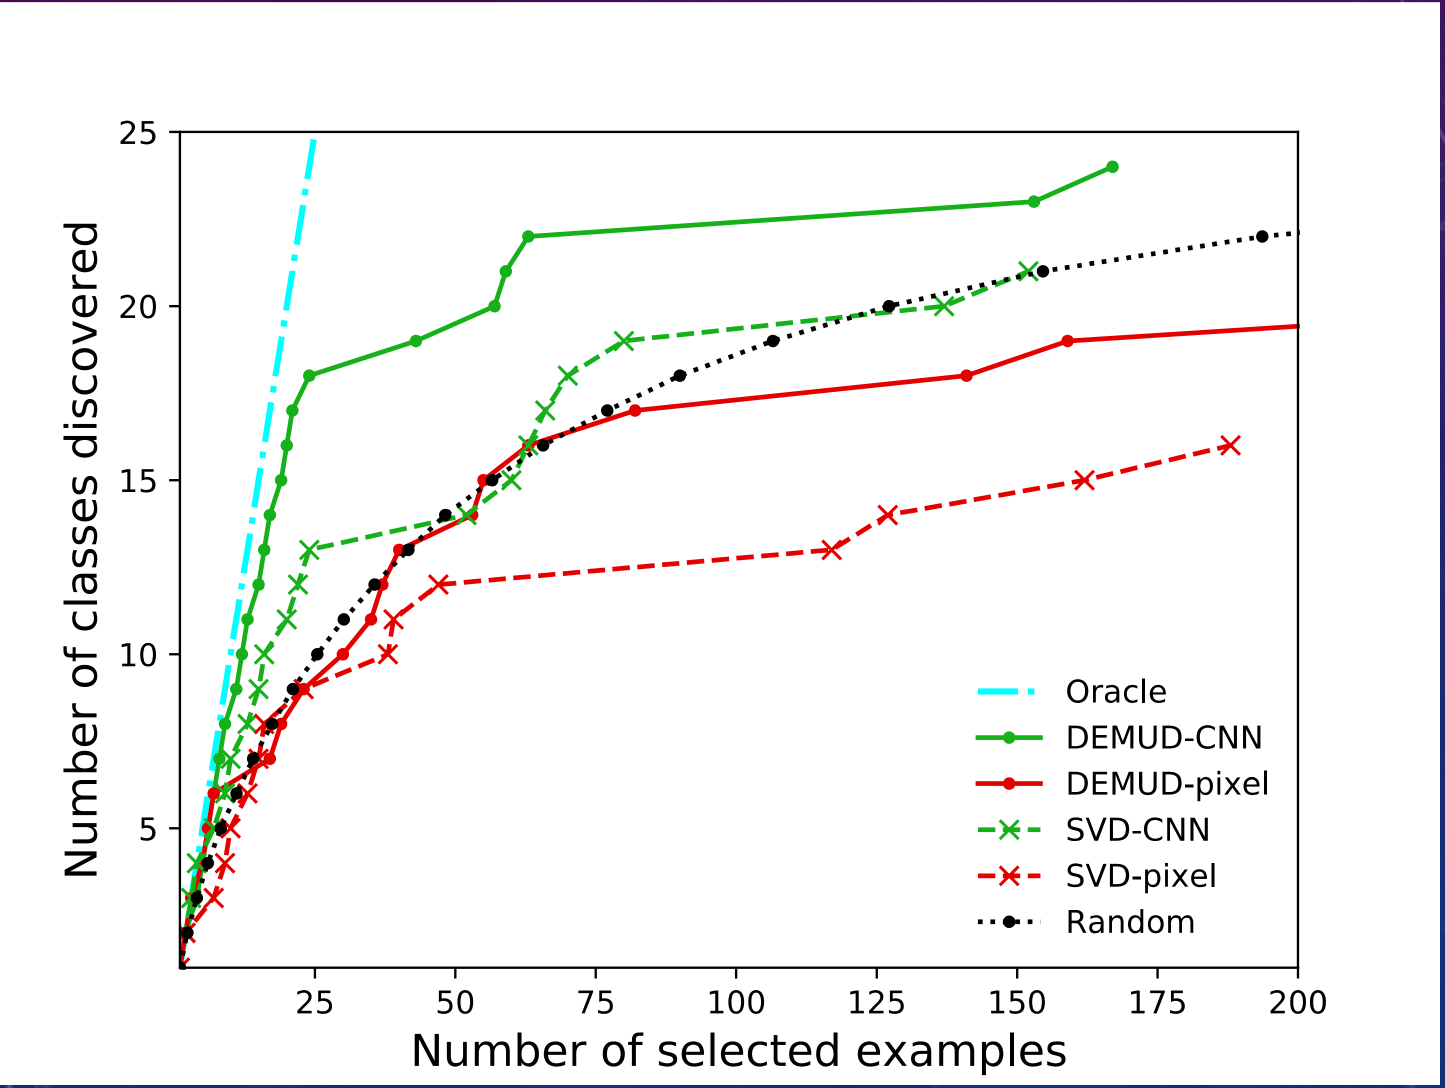

Explanations – MSL Rover images

Experiments – Insect images

- 1362 images: stoneflies from the Pacific Northwest
- 9 classes
- Uneven distribution

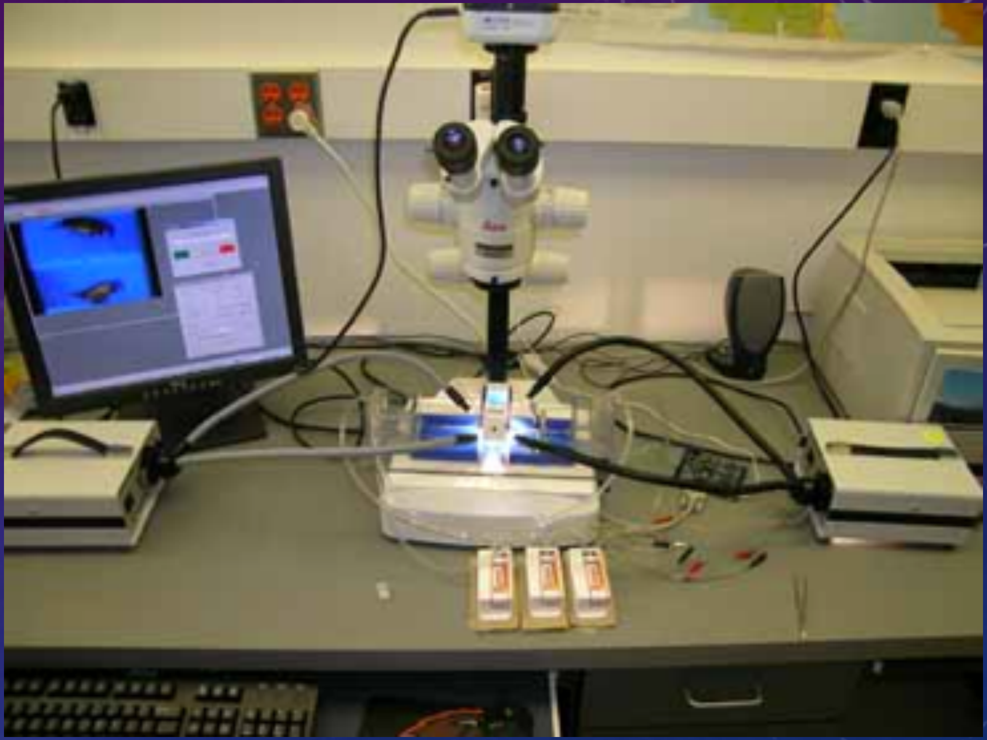

Dietterich et al., Oregon State Univ.

Experiments – Insect images

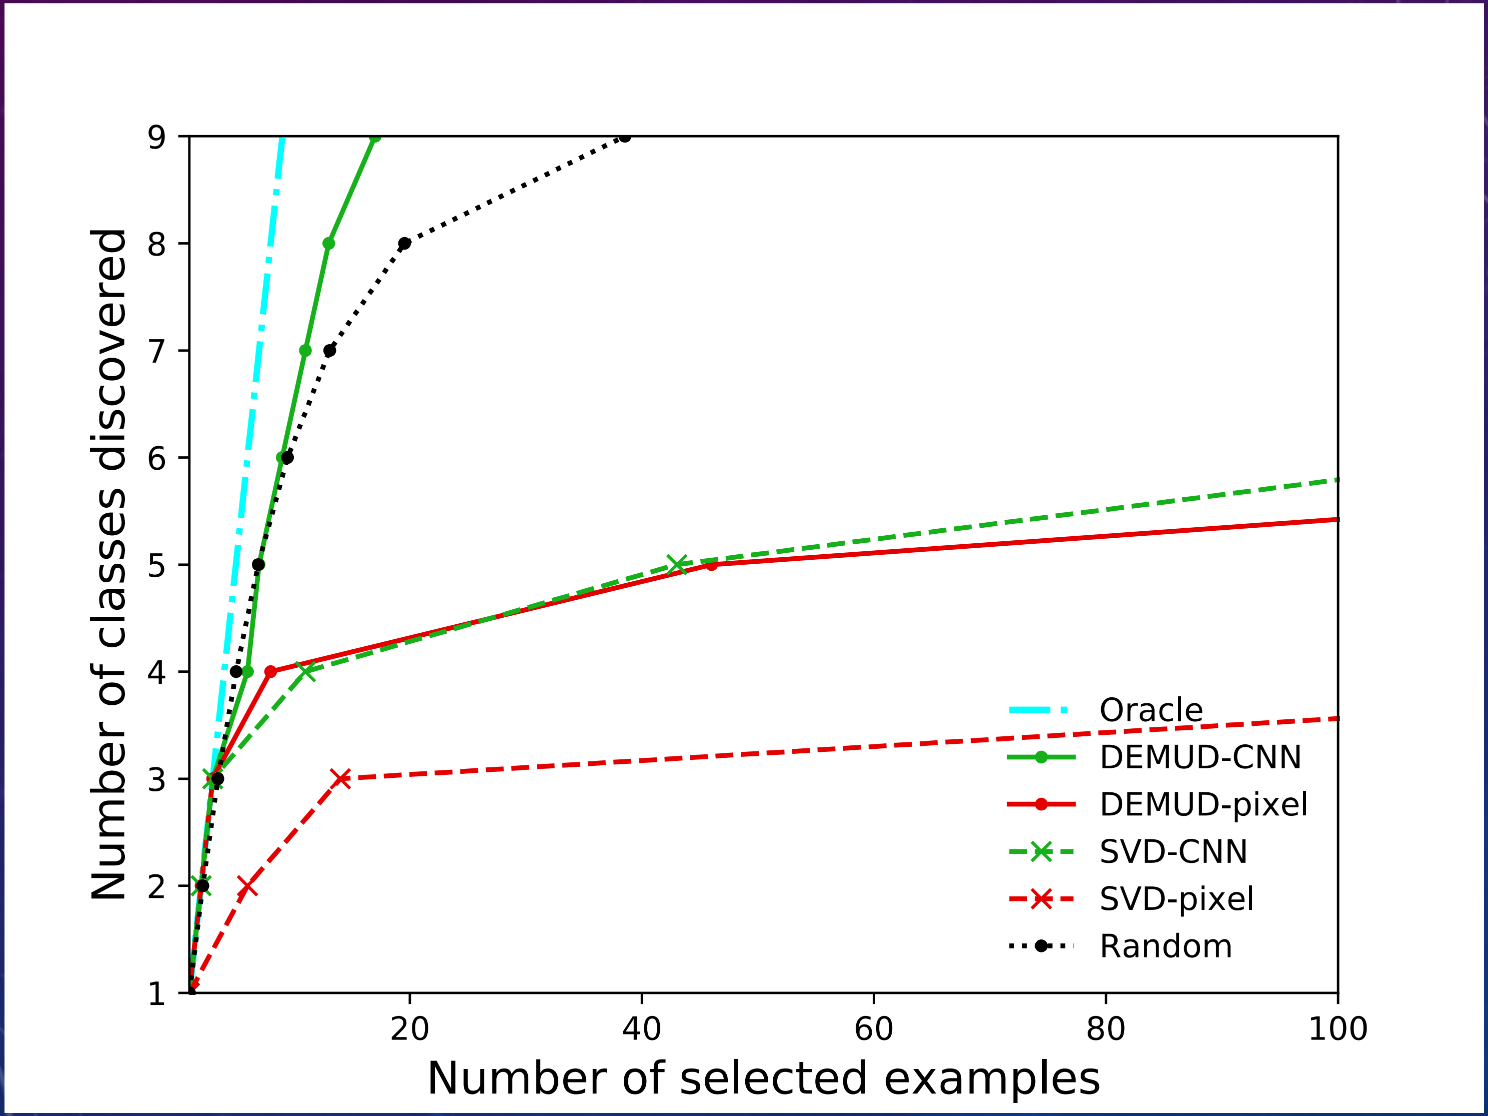

Explanations – Insect images

Image | Simplified image | DEMUD knows | What's new

Zapada

Calineuria californica

Hesperoperla pacifica

Summary

- Machine learning can aid scientific investigations
  - Fast discovery in large or complex data sets
- Interpretable machine learning is vital
  - DEMUD algorithm
    - Quickly discovers new classes
    - Provides explanations

- Examples: ChemCam spectra, Mars rover images, stonefly insect images
- Next: Re-train neural network to specialize on images of interest

Thank you: NASA Planetary Data System (PDS) Imaging Node