Automated Feature Detection

Keck Institute of Space Studies
"Gazing at the Solar System"

Pasadena, CA

2014/06/18

Michael C. Burl
Machine Learning and Instrument Autonomy Group
Jet Propulsion Laboratory, California Institute of Technology
Michael.C.Burl@jpl.nasa.gov

with contributions from many others:
(A. Ansar, B. Bue, T. Fuchs, M. Garay, C. Padgett, D. Thompson, K. Wagstaff, ...)

URS244591

Outline

- Intensity/pattern-based detection of planetary landforms
  - Volcanoes in the Magellan SAR Global Mapping of Venus
    - Cataloging landforms: Planetary Science meets Big Data
    - Rise of the (Learning) Machines
  - Craters
    - Scale invariance
- Exploiting Height for Feature Detection
  - Bue & Stepinski – laser altimetry for Mars crater detection and terrain classification
  - Car detection
- Discovery
  - What if you don't know what you're looking for?
  - What if there are limited examples or target is highly variable in appearance?
  - Treat as an anomaly or salience detection problem
    - Staring allows use of height or temporal change as a signal for salience
- Triggering/Cueing & Event-based Observations
  - Use of onboard analysis to direct future observations
  - Pointing when and where something interesting is happening/about to happen
- Sensor Webs
  - Combining multiple instruments into a sophisticated observational system
  - Value of persistence

Motivation for Automated Image Analysis

Analysis of Large Datasets

Magellan (1989-94)

- SAR Mapping of over 98% of the surface of Venus!
- Returned more data than all previous planetary missions combined
- 30,000 images, hundreds of CD-ROMs
- An estimated 10^6 volcanoes
- An estimated 10 to 20 man-years to catalog manually.

Continued technological improvements in sensors and collection strategies (e.g., **staring**) will yield even larger datasets.

New Science Opportunities Enabled by Onboard Analysis

Detecting and Retargeting Objects of Interest, cueing other sensors.

Downlink Prioritization and Derived Products

Vigilant Monitoring (watching a location for changes)

Closing Feedback Loops Onboard

Magellan SAR Imagery of Venus

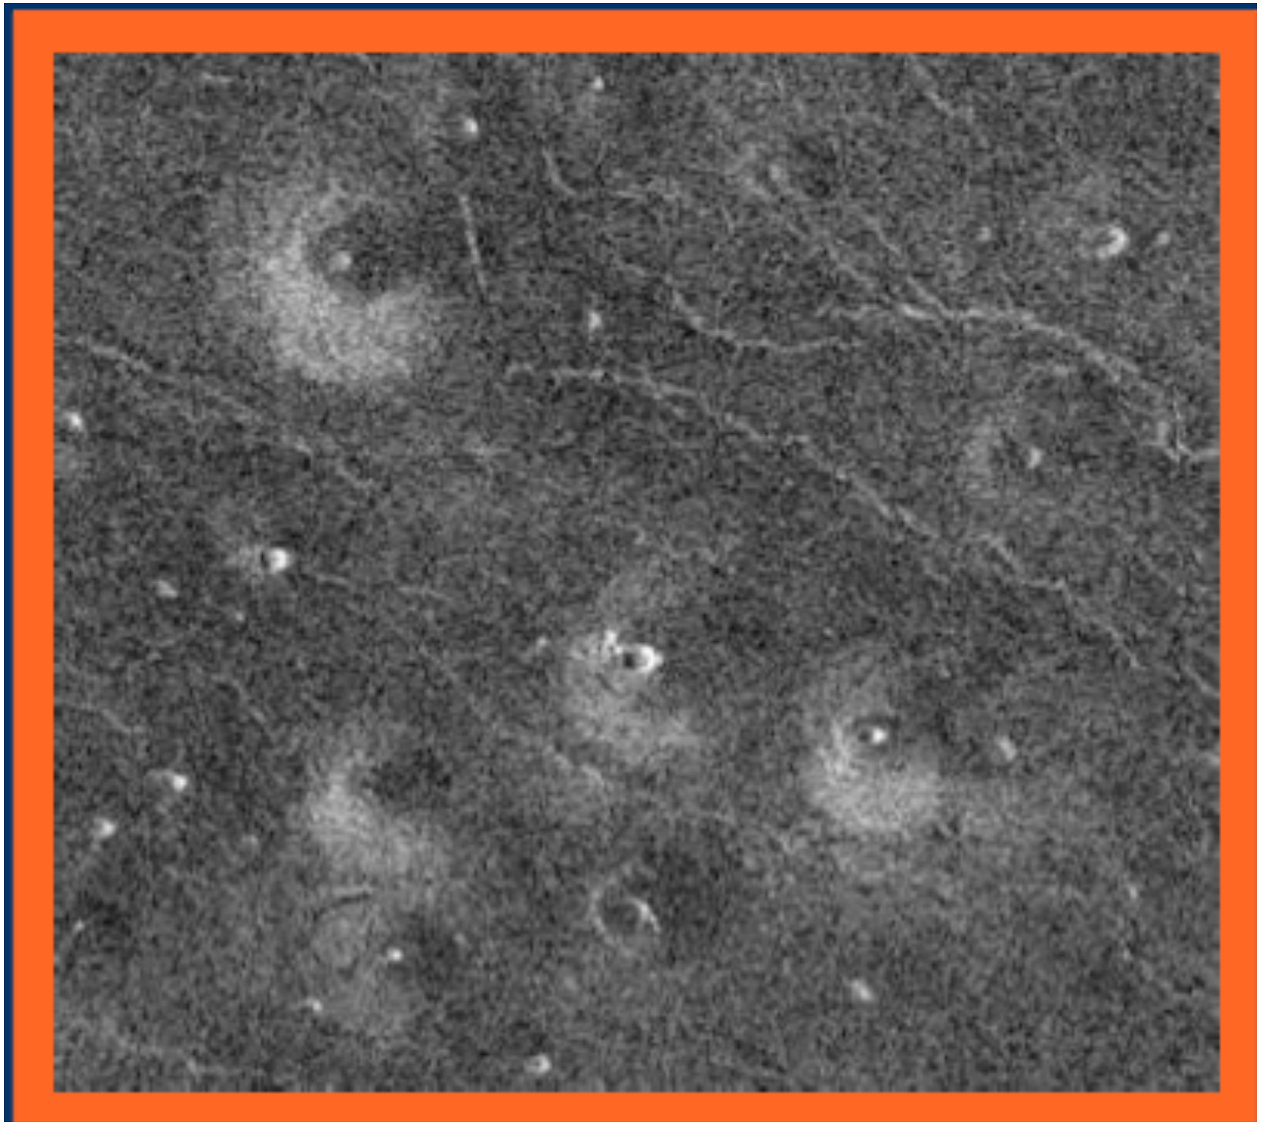

Developed learning-based recognizers that were trained from limited examples labeled by planetary scientists.

Applied recognizer in production mode to generate catalog of volcano locations and attributes (size, subtype, etc.)

Recognition of Geological Objects

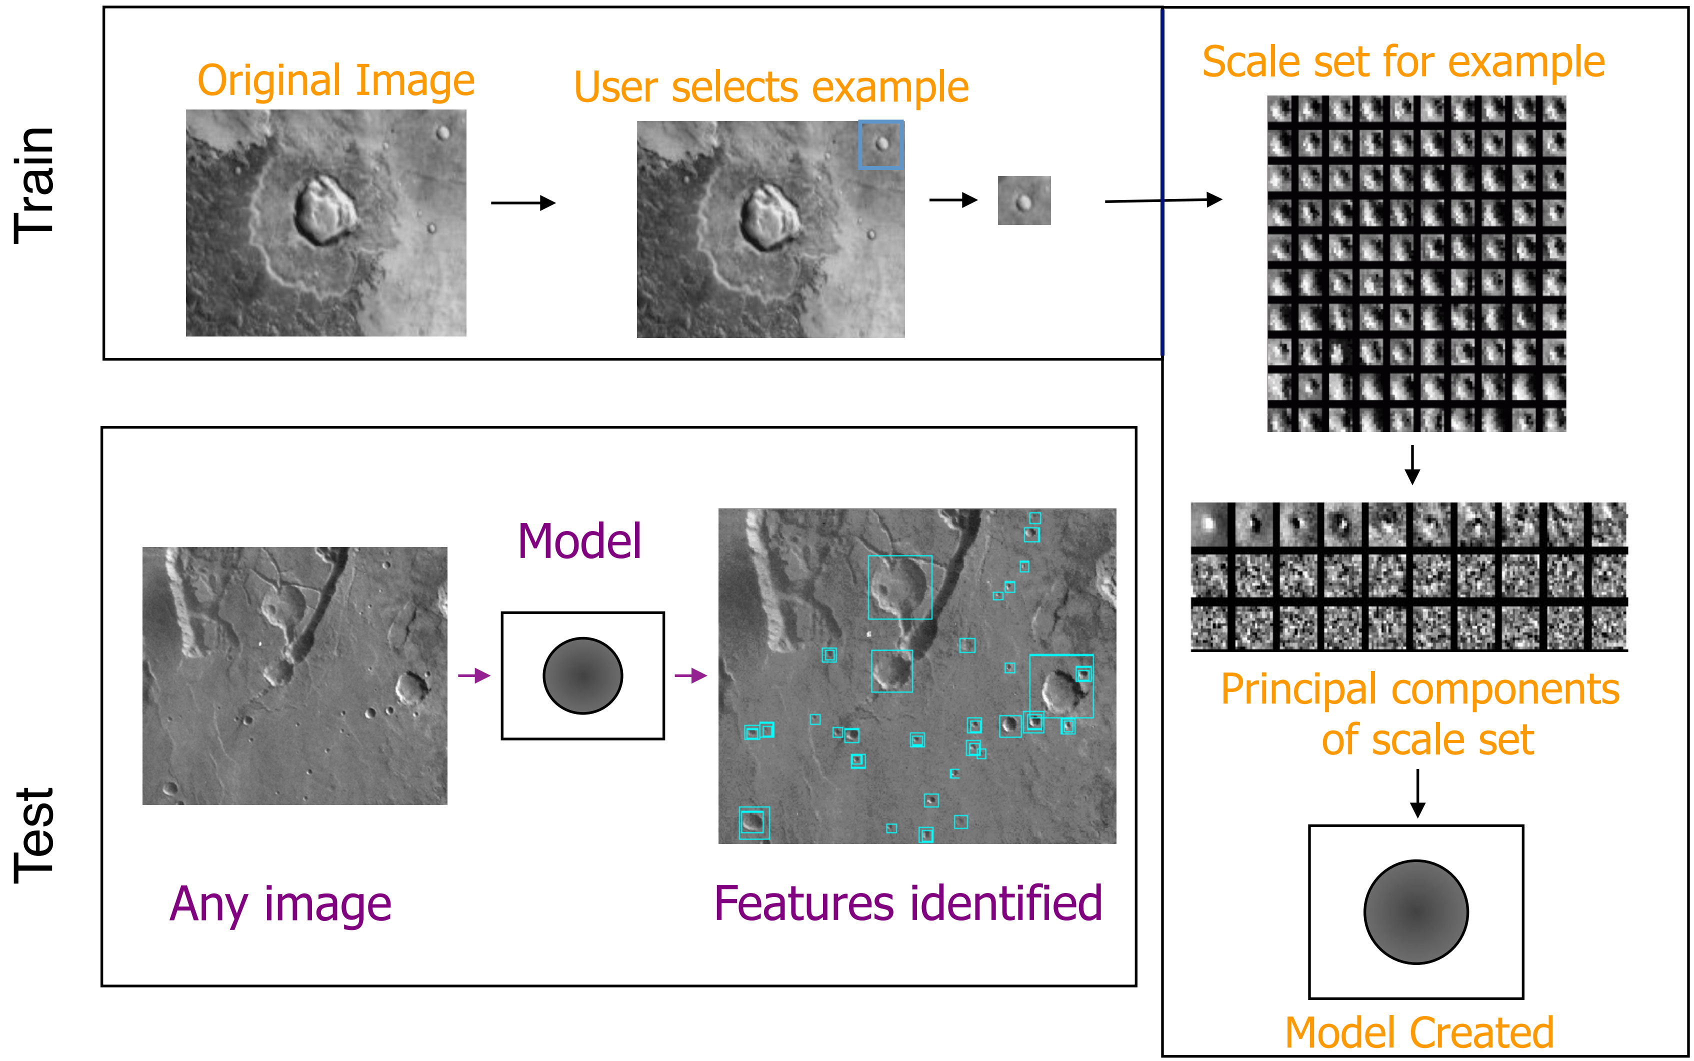

CSVM normalized mi00n242

opps = 301
dets = 197
fas = 72

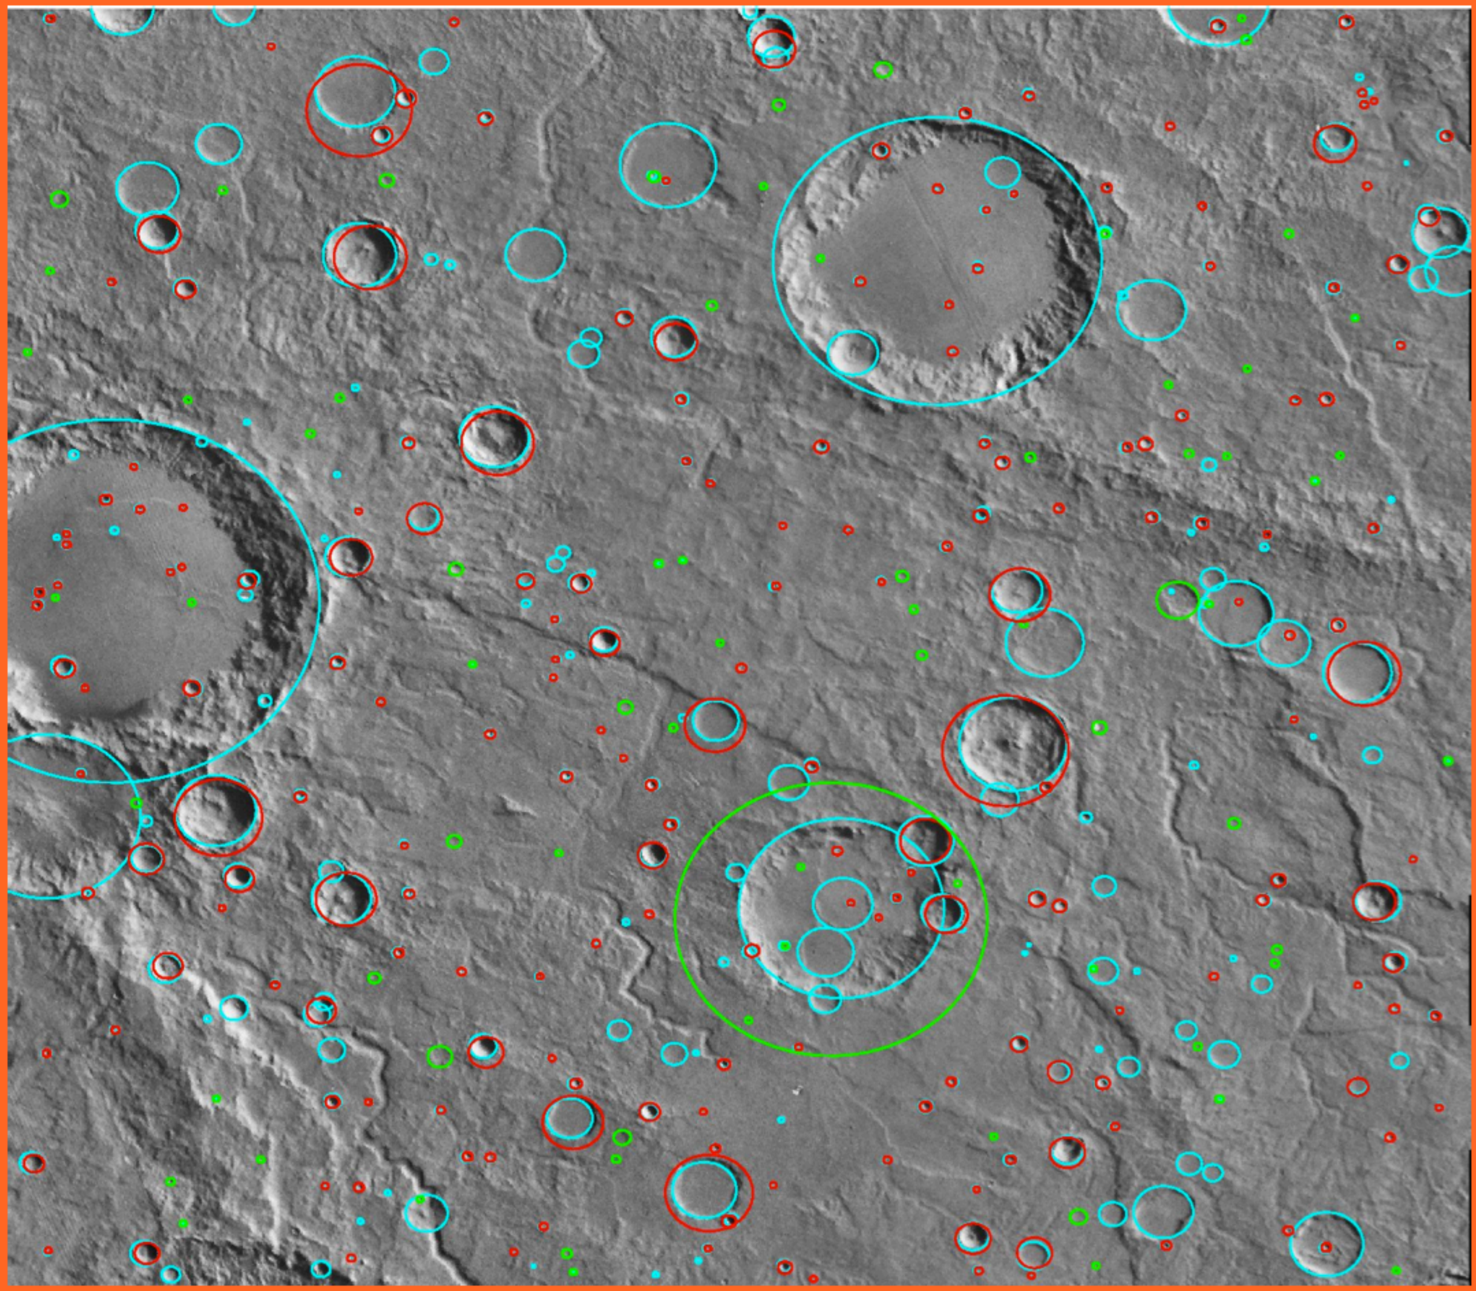

Legend

ground truth
detection
false alarm

Final ROC Curves

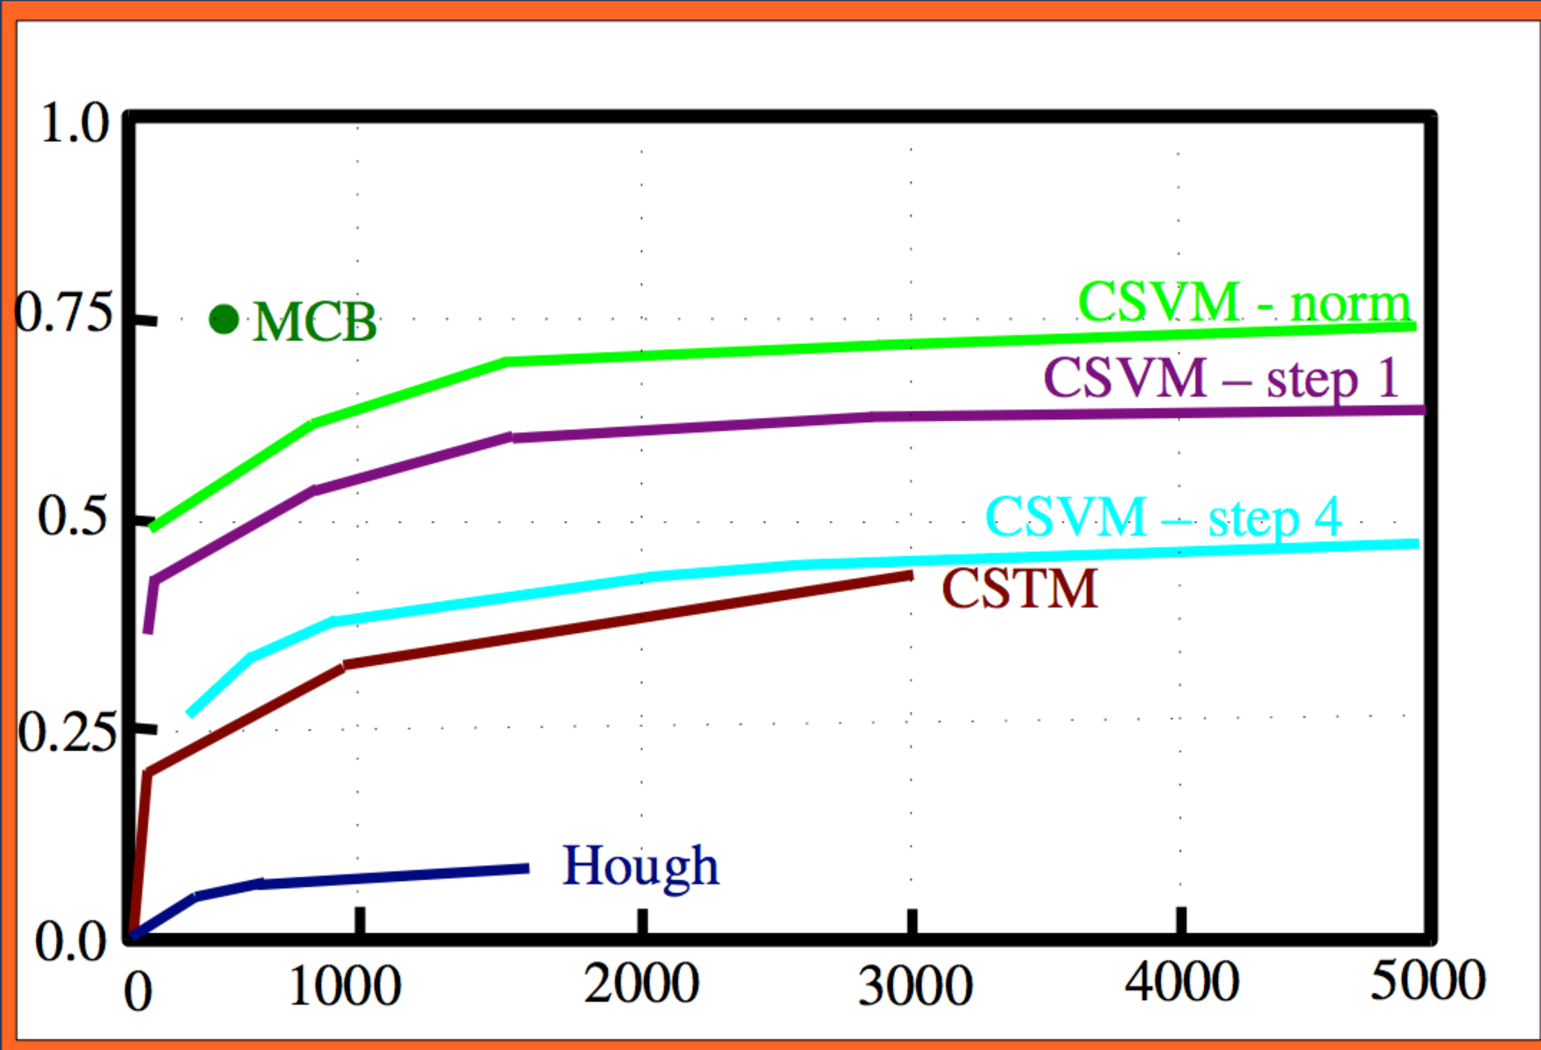

Detections based on Height (MOLA)

(B. Bue, T. Stepinski, et al)

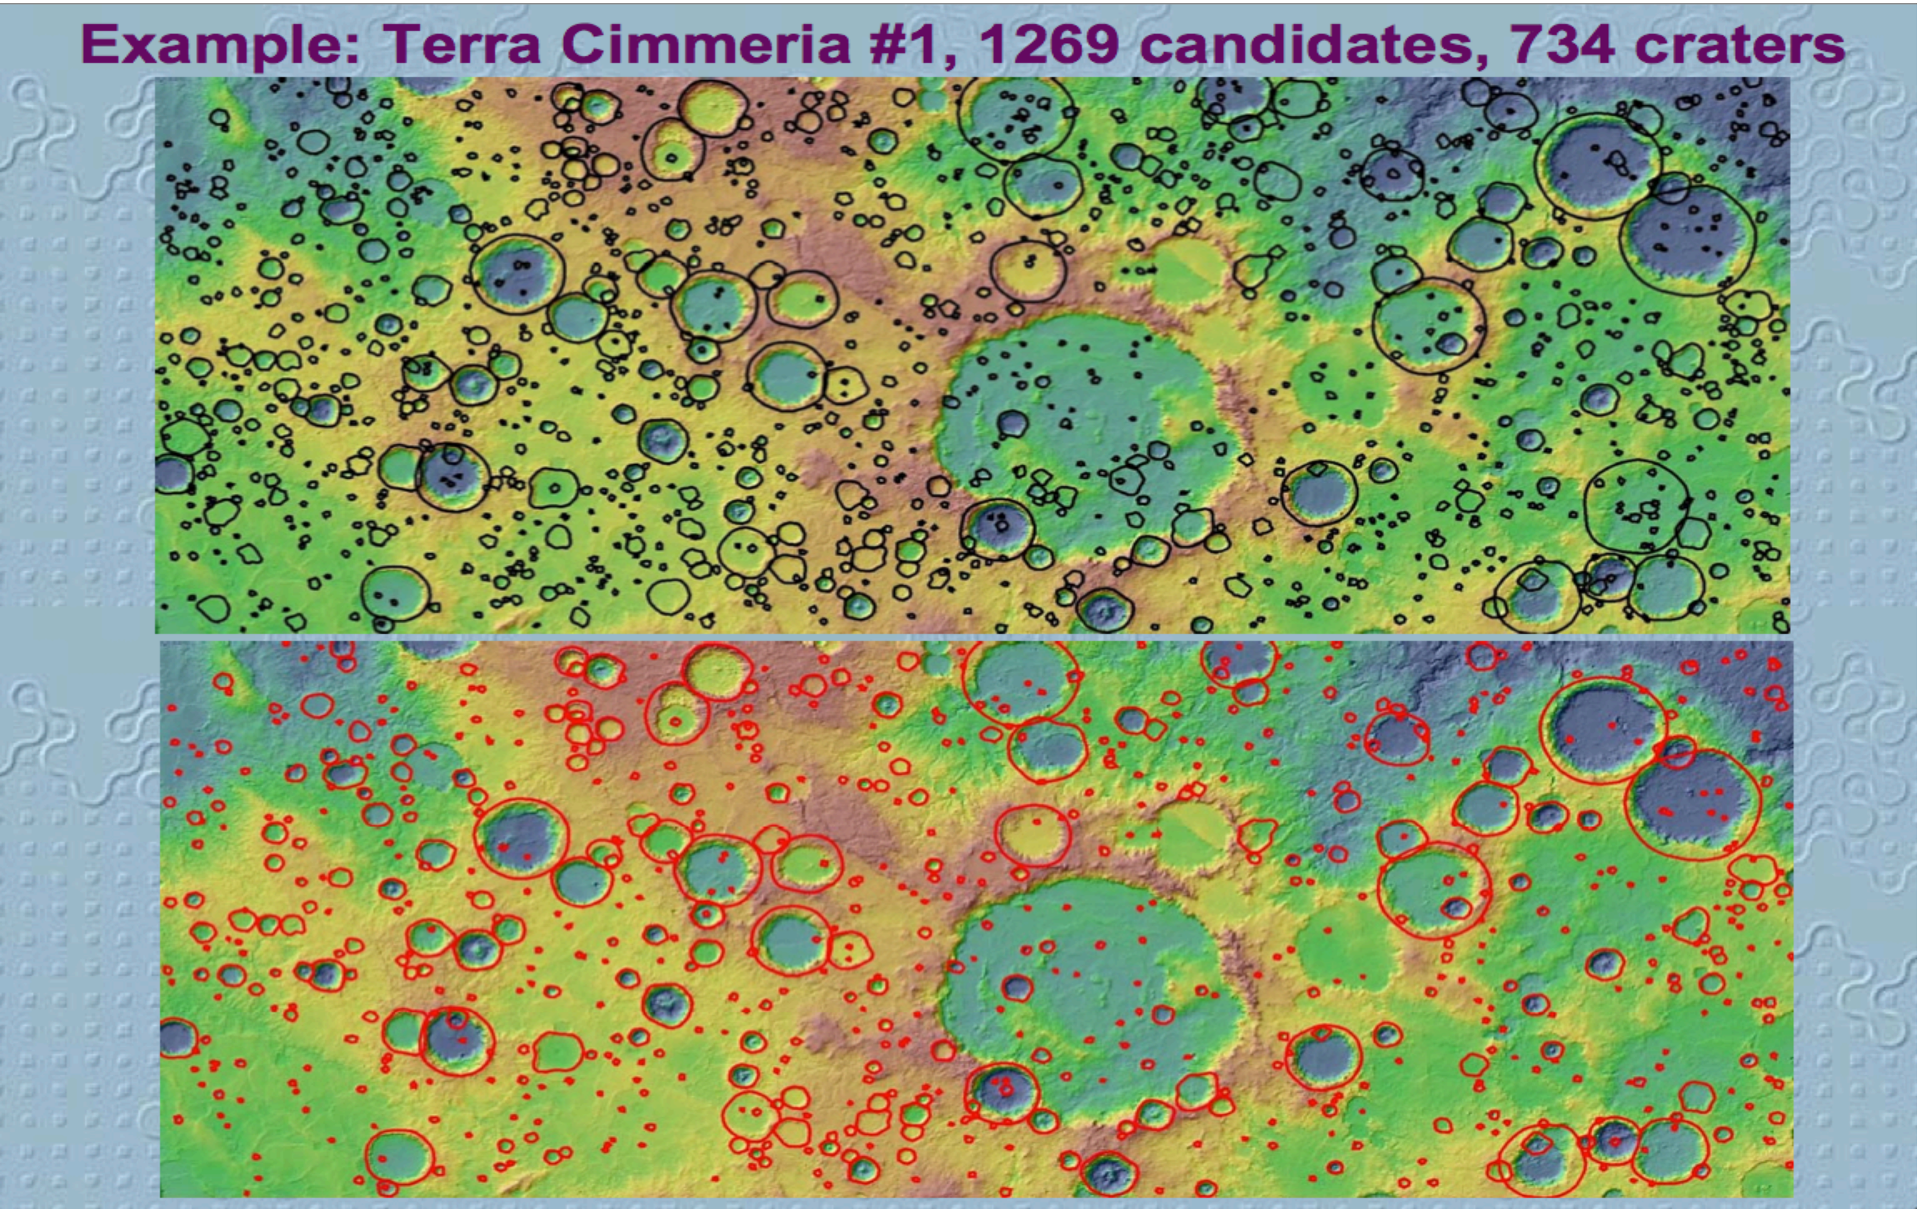

Terra Cimmeria

Terra Cimmeria Classification

(B. Bue, T. Stepinski, et al)

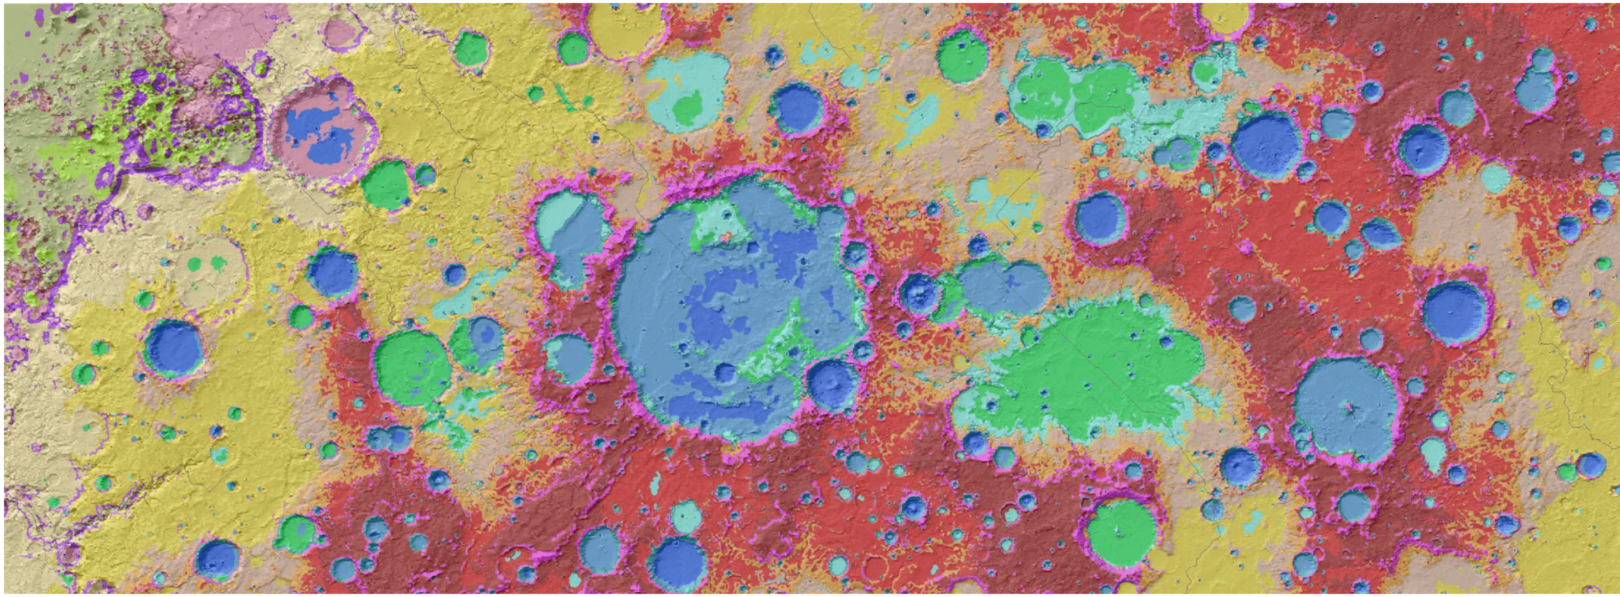

HIGHLANDS PLATEAU

- Inter-crater plateau located at high elevations.
- Inter-crater plateau located at medium-high elevations, small slope.
- Inter-crater plateau located at medium-high elevations, larger slope.
- Inter-crater plateau located at medium elevations.
- Inter-crater plateau located at medium-low elevations.
- Inter-crater plateau located at low elevations, larger slope.

CRATERS

- Terrain inside craters located deep below crater rim.
- Terrain inside craters located at medium depth below crater rim.
- Terrain inside craters located on crater walls, close to the rim.
- Terrain located inside shallow craters and other basins.
- Terrain located at edges of shallow craters and at very shallow basins.

LOWLANDS

- Low, smooth terrain.
- Low, rough terrain.
- Higher terrain in lowlands.
- Knobby hills in lowlands.

RIDGES

- Very steep terrain located on the outside walls of craters and on ridges.
- Steep terrain located on the outside walls of craters and on ridges.
- Escarpment between the highlands and and the lowlands.

CHANNELS

- Terrain that constitutes lower part of major drainage system - channels.

- Studied thoroughly by Irwin and Howard (JGR, vol. 107, No. E7 pp 10-1 – 10-23, 2002)
- 5 classification groupings, highlands, lowlands, craters, ridges & channels.
- Classification recognizes highland-lowland border.

- Region area: 10^6 km^2
- 593 km across at the equator
- 5,303,888 classifiable pixels
- 6D Digital Topography Model [Bue and Stepinski, Computers & Geosciences, 2006]

Detecting Stationary Cars in Wide Area Motion Imagery (WAMI) Data
Scored Detections

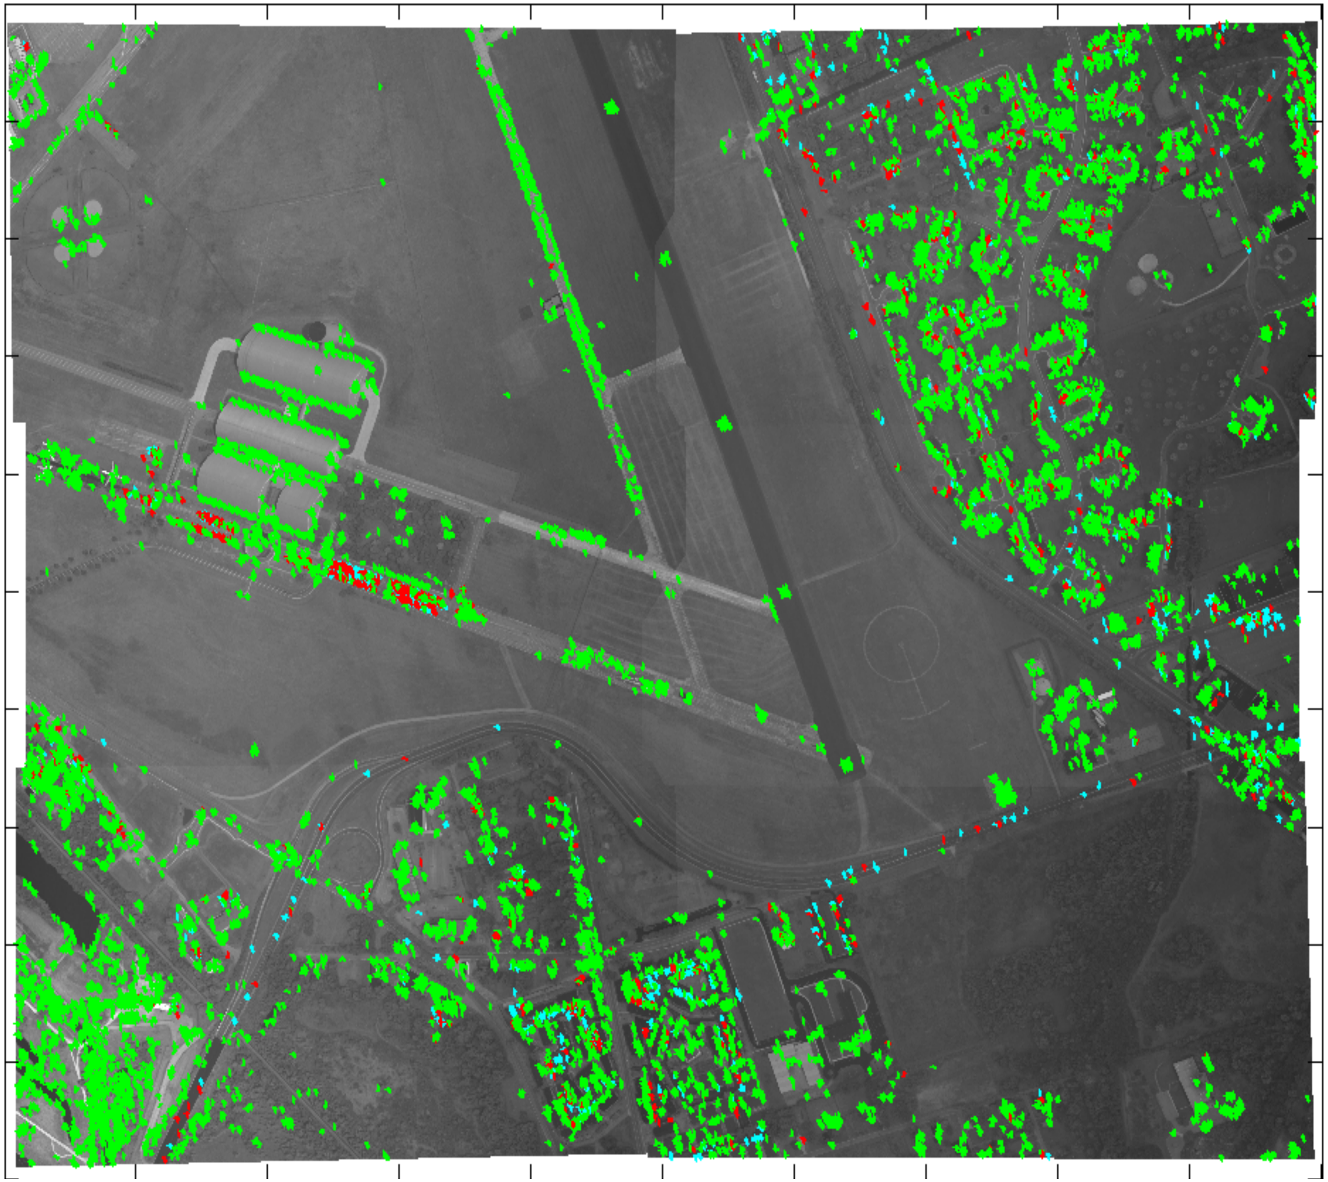

Red=true detection (TP)
Cyan = miss (FN)
Green = false alarm (FP)

Simple detector based on intensity pattern yields large number of false alarms!

Scored Result – Detail1

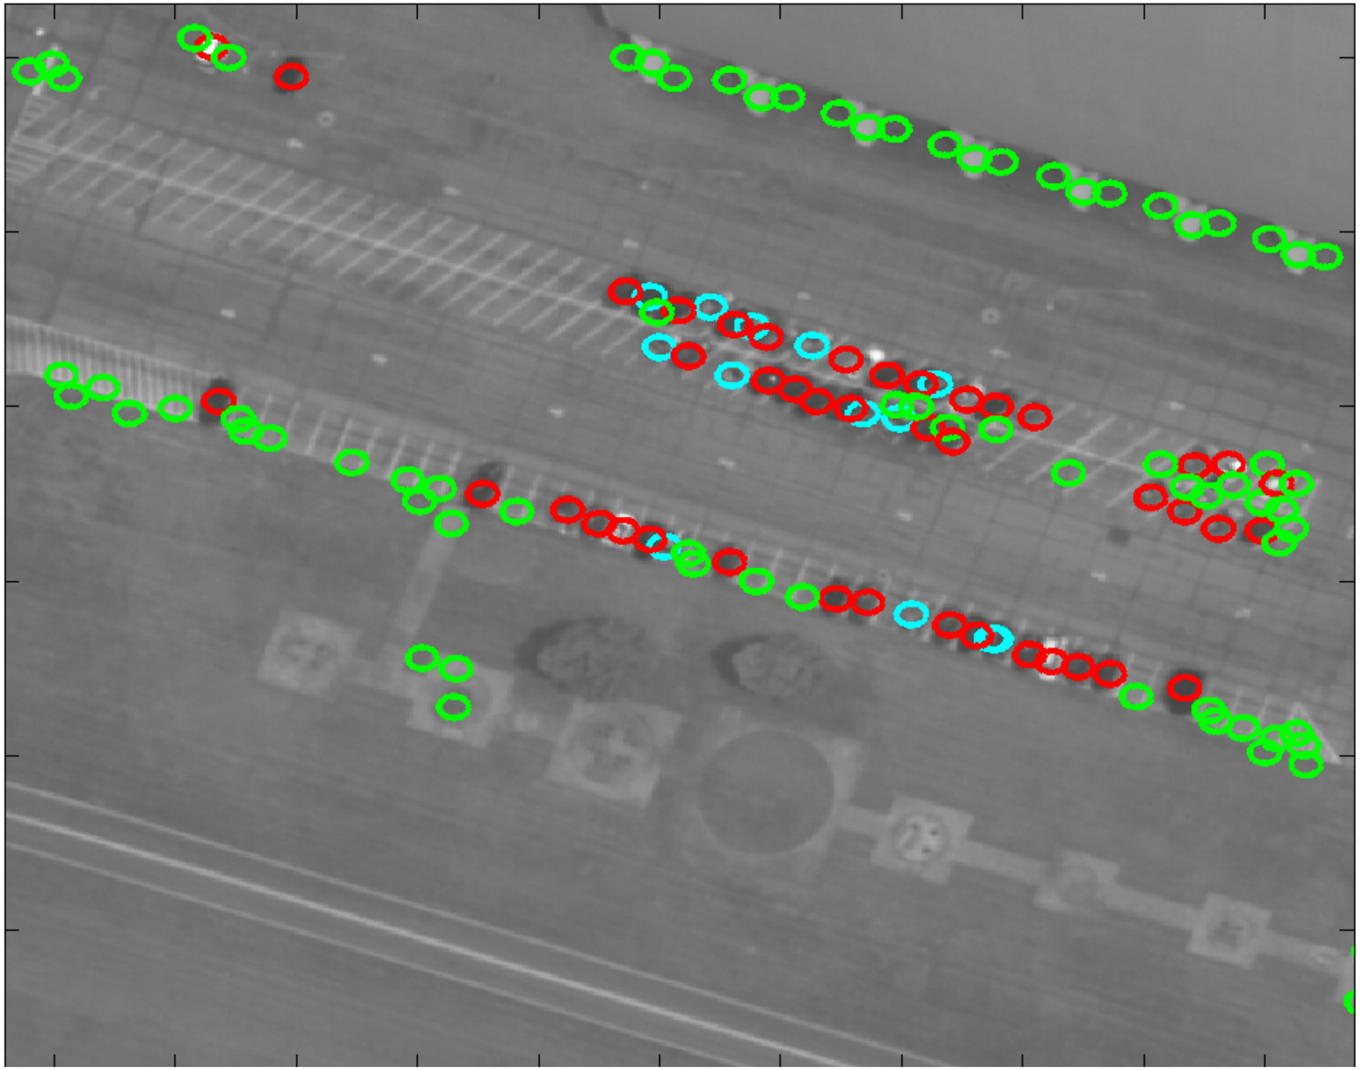

Red=detection (TP); Cyan = miss (FN); Green = false alarm (FP)

Positive Examples

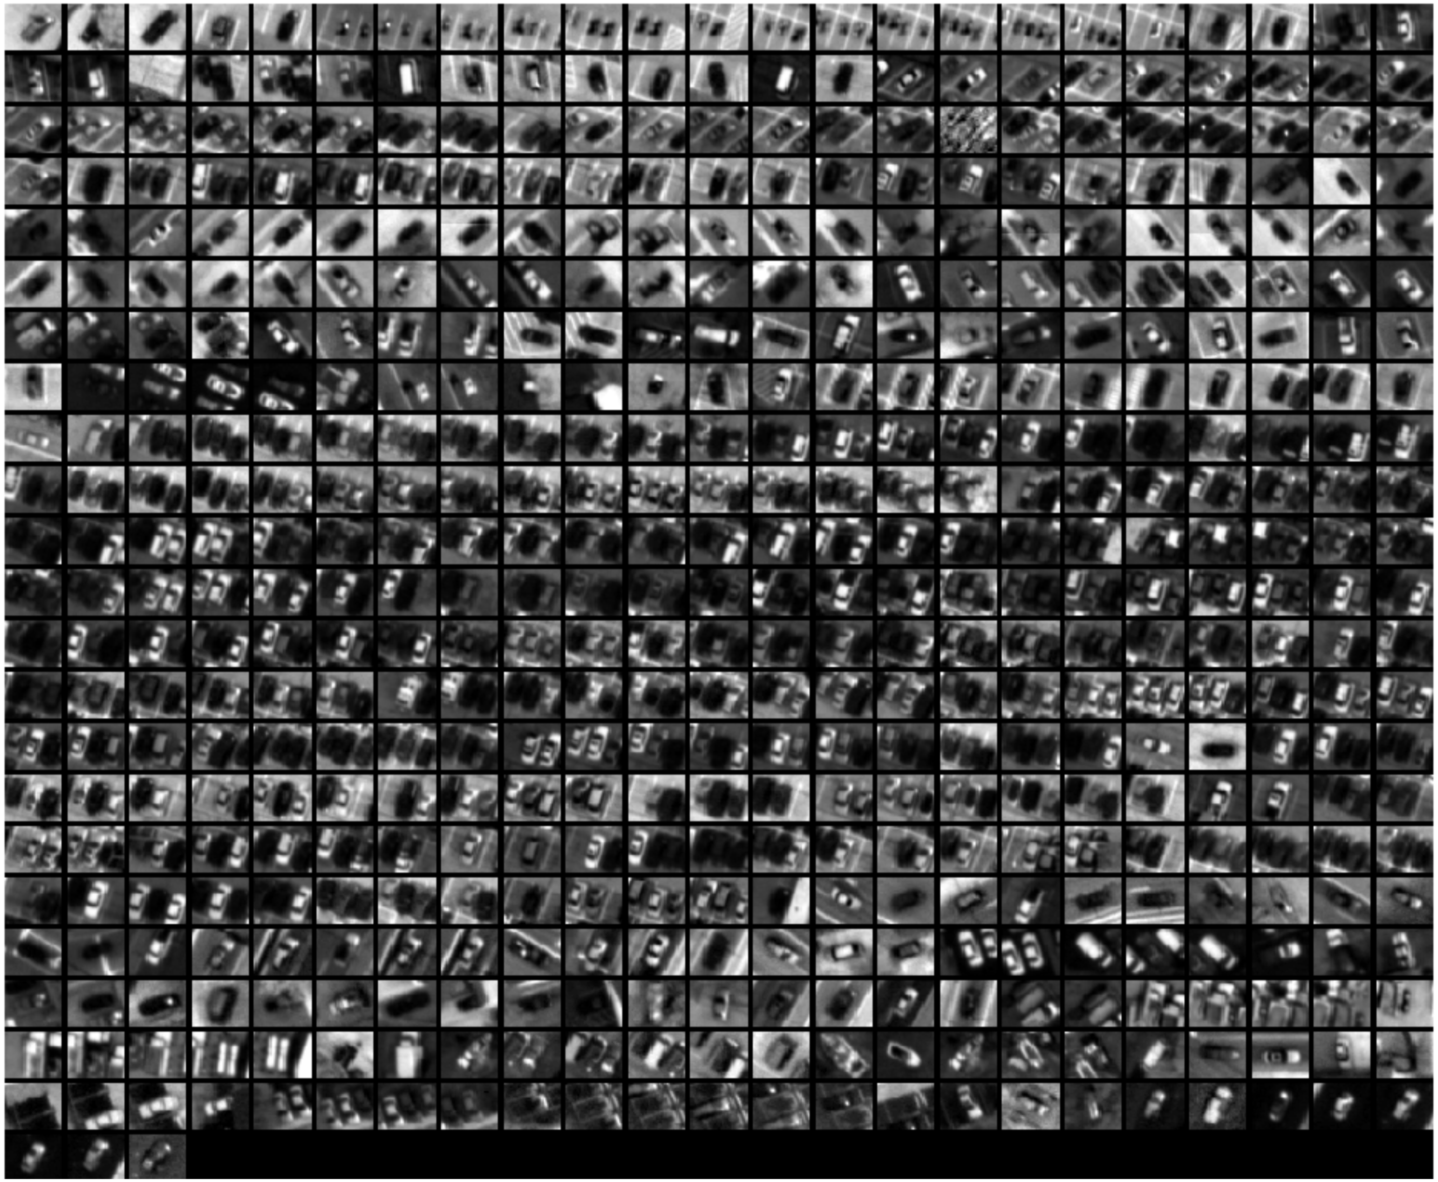

Positive Examples from left 1/3 of image.

33x33 pixel patches

Typical vehicle:
9-12 pixels width
18-27 pixels length

False Alarms (1:1000)

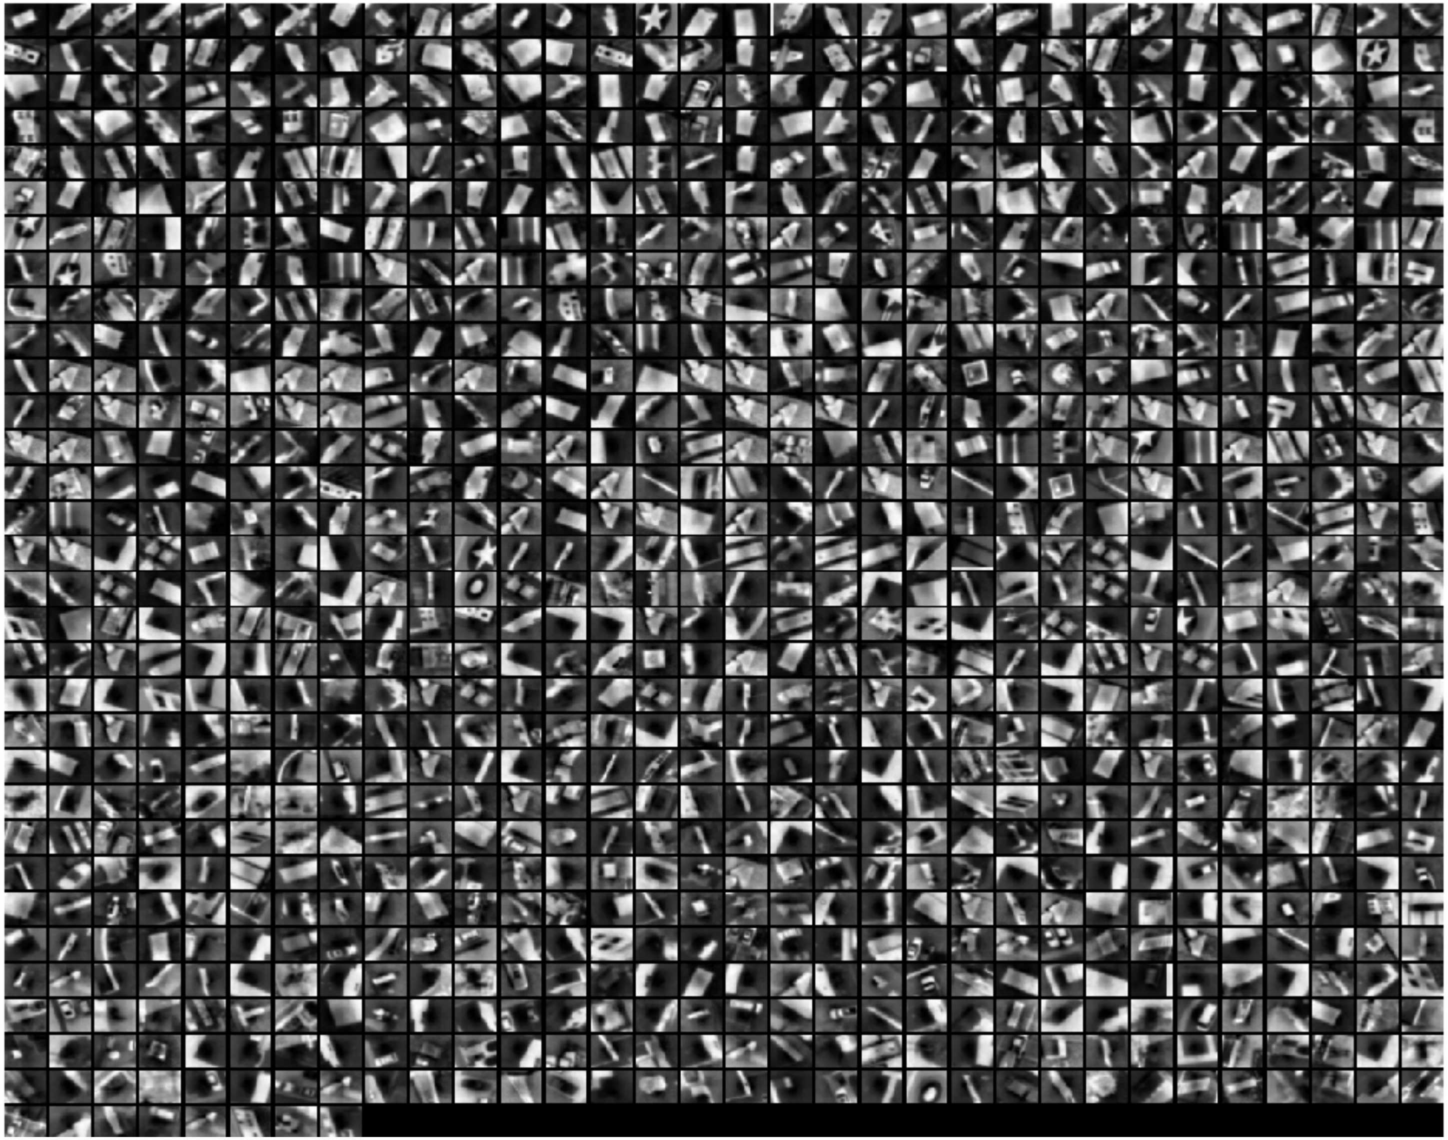

Idea: Use DEM recovery and object height to reject many of these false alarms.

DEM from Multibaseline Stereo

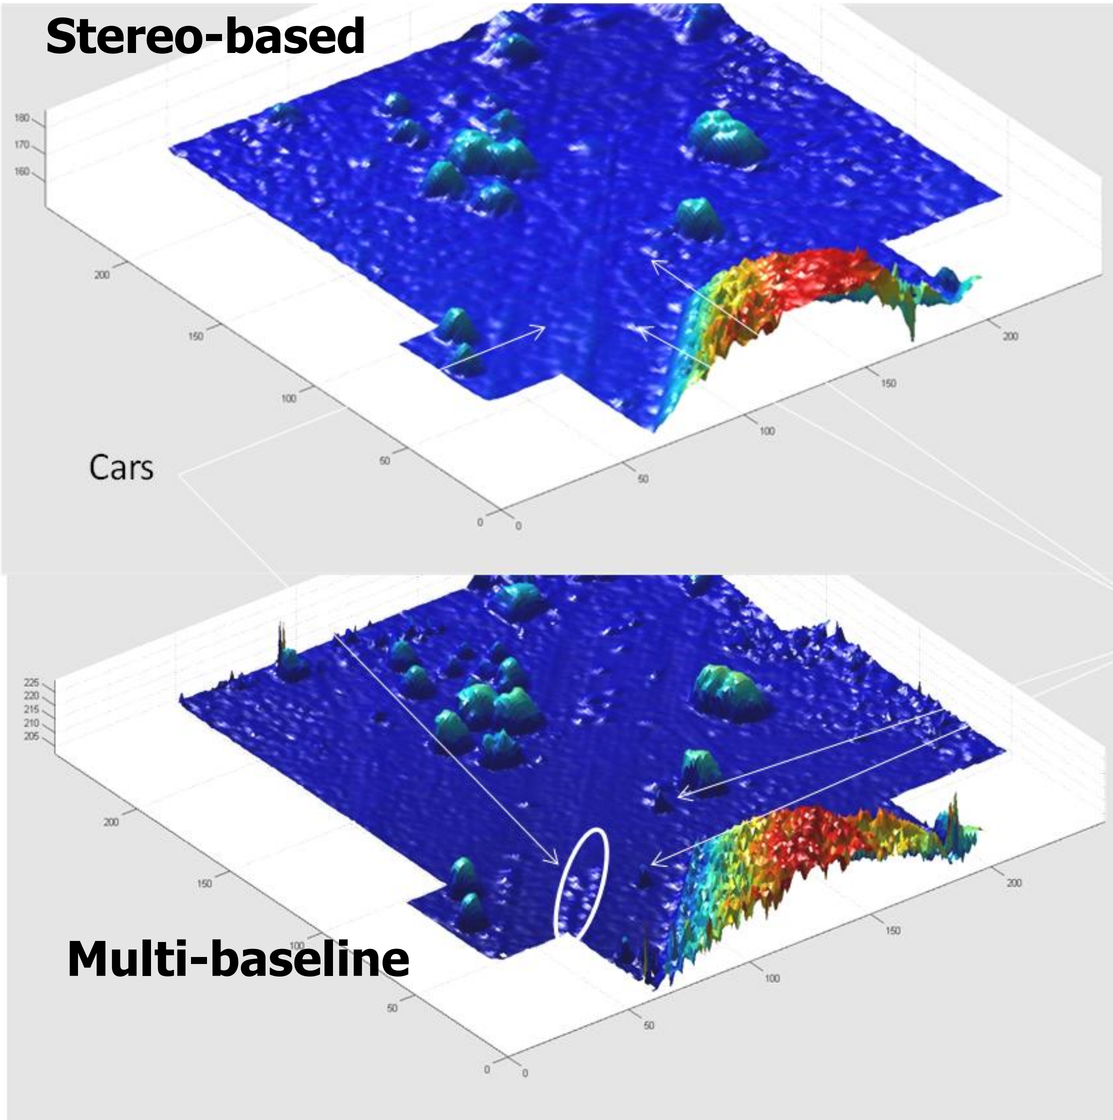

Improved Image-Based, automated, 3D generation
– top picture, stereo based structure from motion
– bottom picture, multi-base line structure from motion (more discrimination closer to the ground)

The multi-baseline technique provides better height estimation (over a specified range) and spatial resolution. Lamp posts and cars can be picked out in the bottom image but not in the top

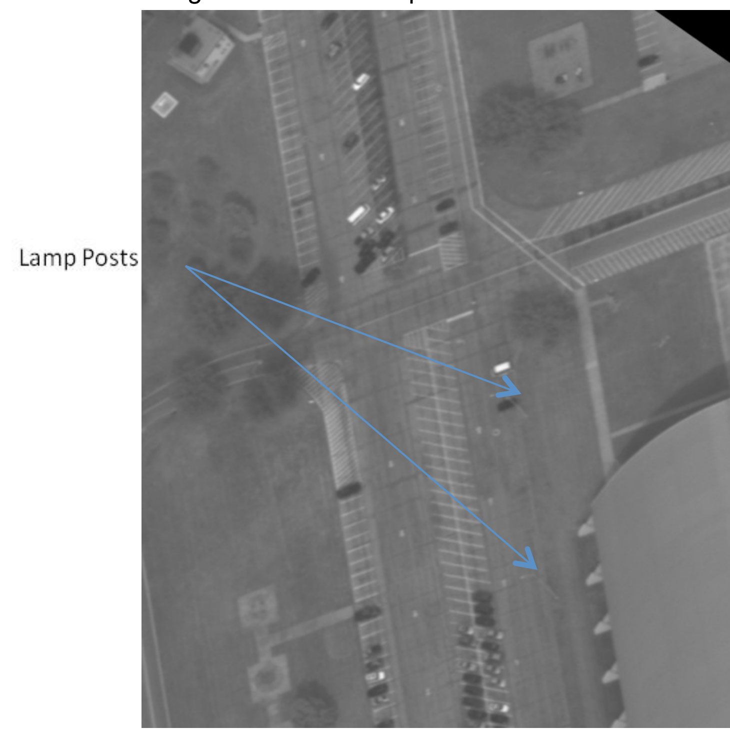

(A. Ansar, C. Padgett, et al)

3D using Multi-Baseline Stereo

(A. Ansar, C. Padgett, et al)

Image chips at ½ native resolution (210mm lenses)

Multi-baseline stereo algorithm rectifies arbitrarily many images to plane slices parallel to ground and picks best slice for each pixel

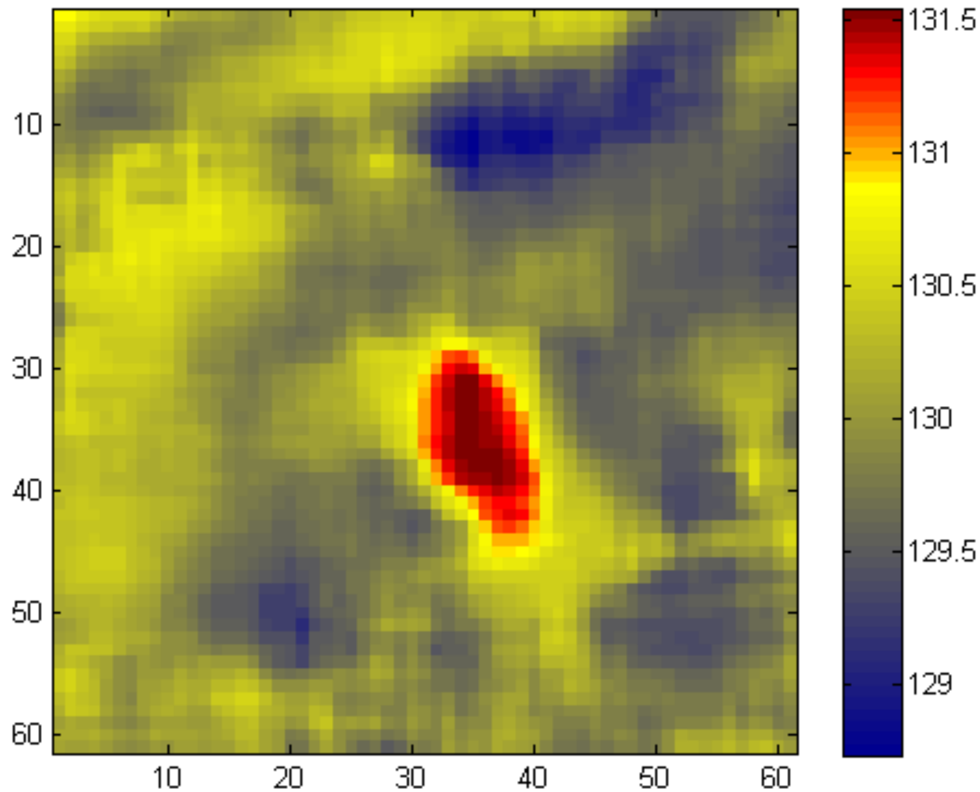

Elevation map for cropped region around car. Area around car ~1.5 m higher than neighboring ground plane. Higher image resolution might address some remaining noise issues.

3D mesh data at car derived using multi-baseline stereo

(A. Ansar, C. Padgett, et al)

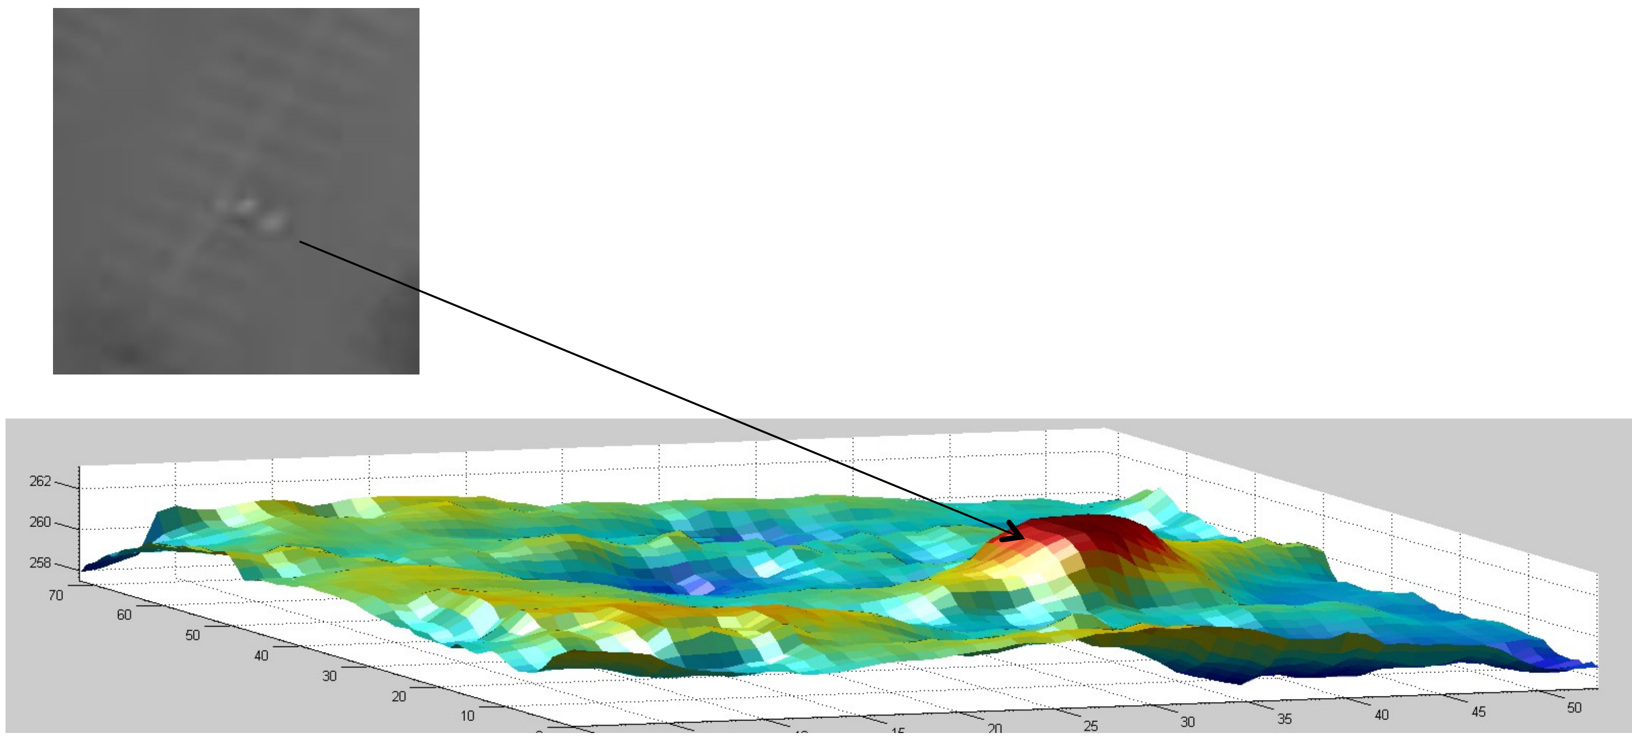

Incorporation of a “height test” reduces false alarms 15x, but loses a significant number of cars, as well. Have to be careful!

Taxonomy of Algorithms for Finding Objects

Low ← **A Priori Knowledge** → *High*

- No training examples
- No idea
- No regularity

Discovery

- One example
- Conceptual notion

Queries

Volcanoes, Craters, Cars

- Many training examples
- Regularity in appearance
- Sometimes an accurate 3D model is available

Learned Recognizers

Develop fairly precise model of how target objects look and variations thereon.

Taxonomy of Algorithms for Finding Objects

Low ⟵ **A Priori Knowledge** ⟶ *High*

- No training examples
- No idea
- No regularity

Discovery

- One example
- Conceptual notion

Queries

- Many training examples
- Regularity in appearance
- Sometimes an accurate 3D model is available

Learned Recognizers

Detect areas that are anomalous relative to surrounding context. With multi-angle imaging, the height or temporal behavior (e.g., change detection) could be used as an aid for "discovery".

Example: Martian Dunes – Fairly Irregular in Shape/Appearance

Some objects, such as dunes, *may* be easier to detect with a "discovery" type algorithm, rather than with a recognizer that uses a precise model learned from examples.

Idea: Look for places that are different (anomalous) relative to their surrounding context.

Prototype discovery algorithm: compare vector of oriented "early vision" filter responses in an area to responses in neighboring spatial areas.

Evaluation of a Filter-based Discovery Algorithm

Applied prototype to a variety of different datasets without telling it specifically what to look for.

TEST DATA

Imager	Target	Features
Ground–based ESO telescope	Moon	Craters
Magellan SAR	Venus	Volcanoes
Global Surveyor	Mars	Sand dunes
Voyager II	Triton	Ice geysers

Prototype successfully "re–discovers" features known to be of interest to planetary scientists.

V–0258
M.C. Burl

An Alternative Approach to Discovery: Landmark Detection via Contextually-Salient Histograms

(K. Wagstaff, et al)

<u>Landmark</u>: a visually salient region within an image (e.g., crater, dust devil track)
<u>Salience</u>: context-sensitive (dynamic) judgment about how distinctive a region is

Approach: compute statistical difference between pixel and surrounding context

<u>Landmark</u>: contour around high-salience region

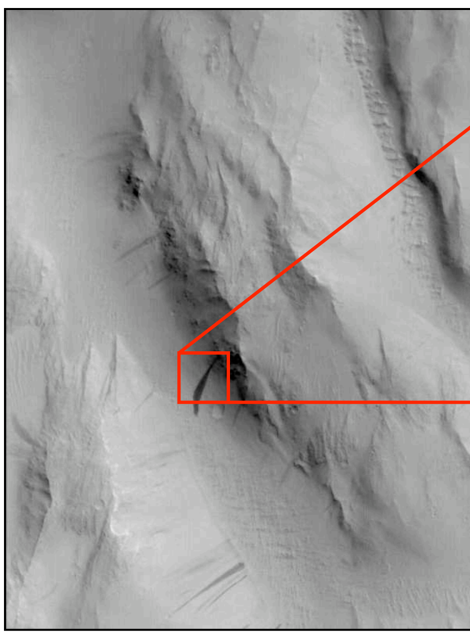

MOC, June 2000

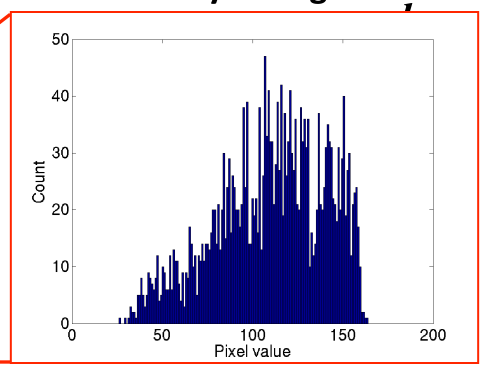

$$S(x,y) = \frac{1}{M}\sum_i |p_{x,y} - i| h_i$$

Salience Map (win=50x50)

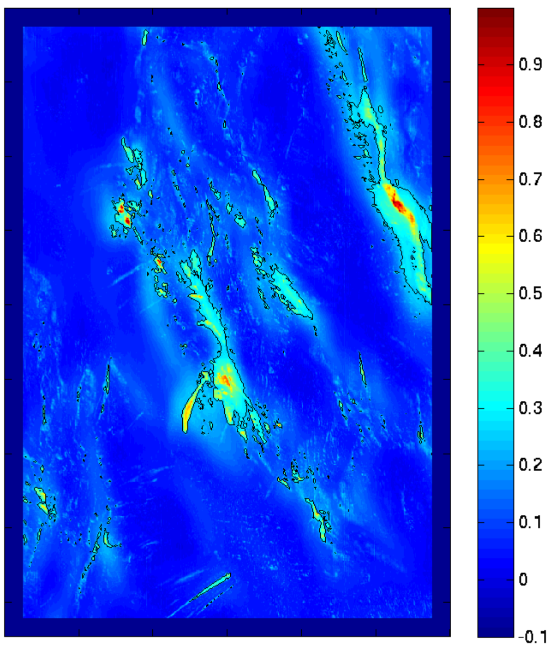

K. Wagstaff, B. Bornstein, L. Mandrake, N. Schorghofer, A. Smith

Triggering/Cueing and Event-based Observations

Rockster – Rock Segmentation through Edge Regrouping

Incoming Image

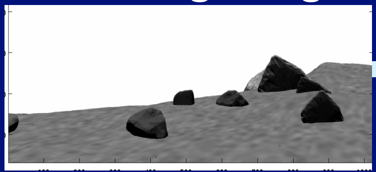

Sky – Ground Segmentation

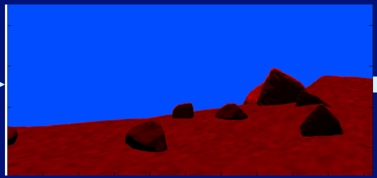

Edge Detection and Linking

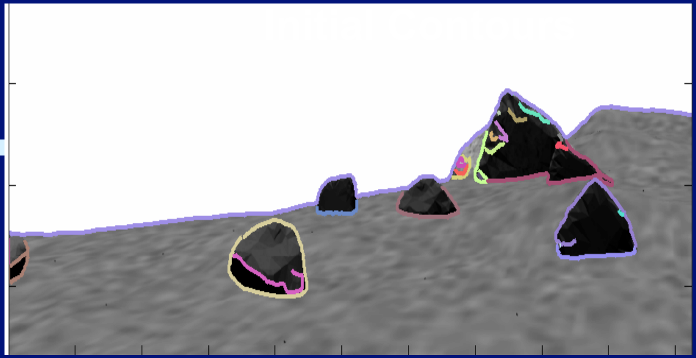

Salient points; splitting; T-junction and gap filling

Revised Contours

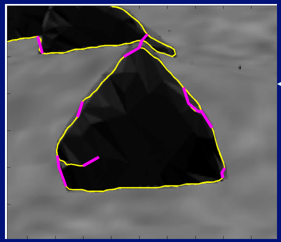

Background flooding

Rock Mask

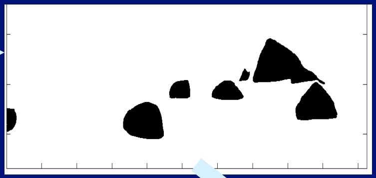

Final Result

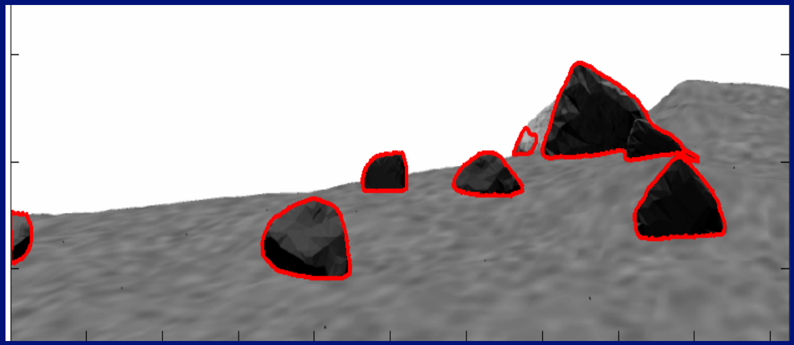

Running onboard MER Opportunity rover as part of AEGIS – enables autonomous targeted follow-up observations with no ground intervention required.

(D. Thompson, et al)

Onboard Analysis for Event-based Observation: Asteroid Flyby

- Targets have diverse morphologies, compositions
- Target locations are not known in advance
- Closest approach may pass quickly (sub-hour timescales)
- Geometry and illumination constraints
- Features of interest are highly localized

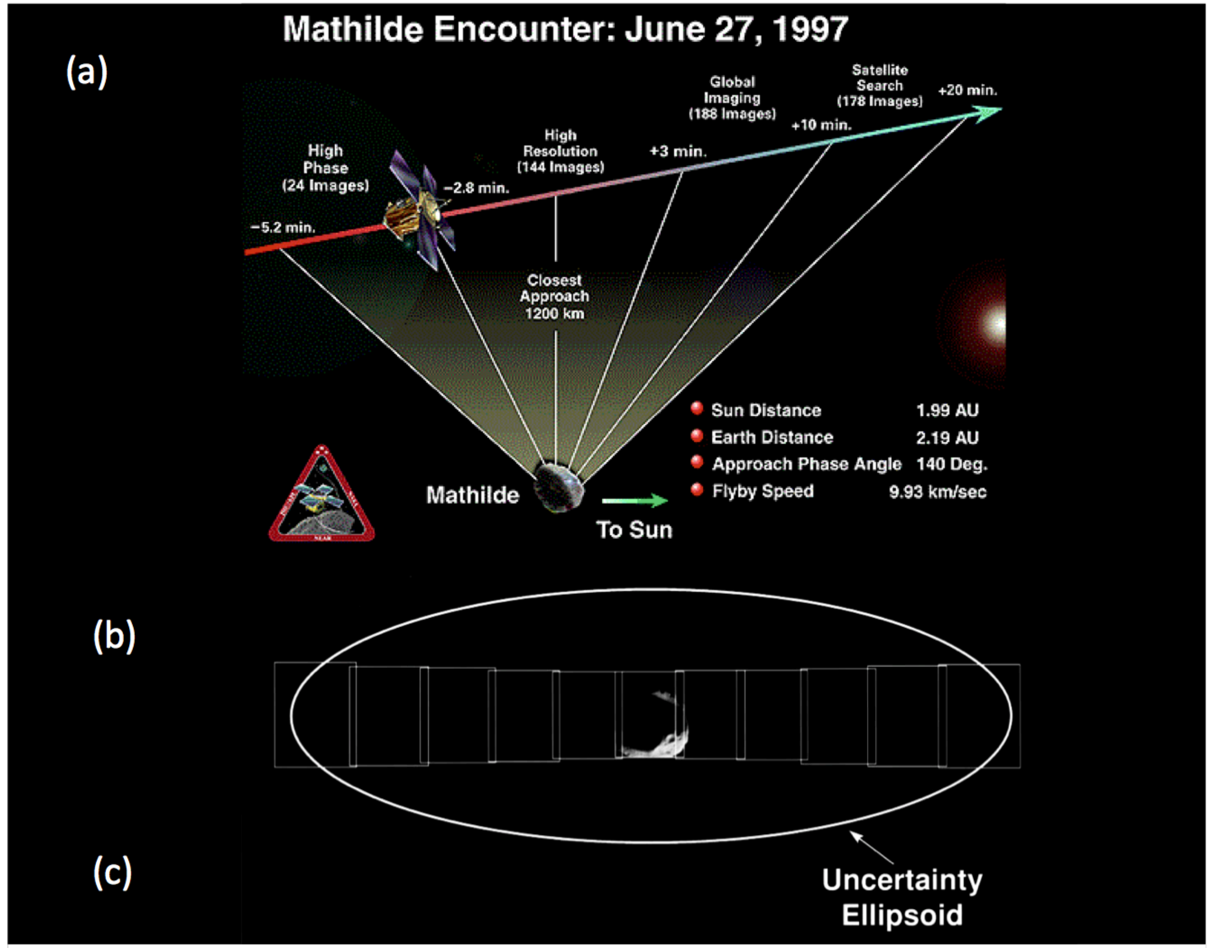

J. Veverka and 16 co authors. Near's flyby of 253 mathilde: Images of a C asteroid. Science, 278:2109, 1997.

(D. Thompson, et al)

Status Quo: respond in days

Processing and analysis

Replanning and sequencing

Onboard analysis: respond in minutes

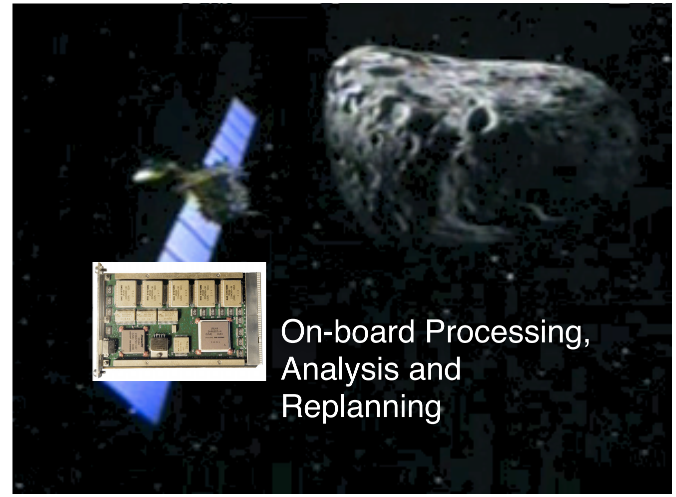

Rosetta graphic courtesy NASA / ESA / US. Rosetta
DSN image courtesy NASA / Caltech / JPL

(D. Thompson, et al)

Plume detection algorithm

Edge detection

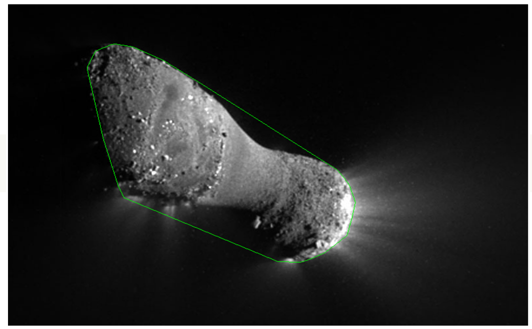

Convex hull

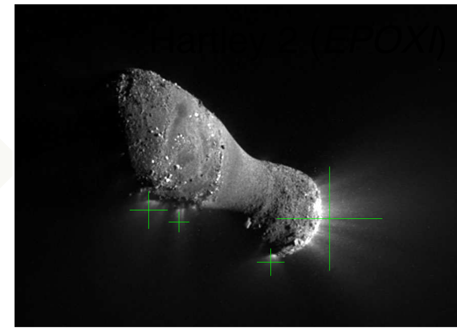

Segmentation and thresholding

- Detect bright material beyond the limb
- Enable monitoring campaigns, target-relative data acquisition
- Detects most plumes with zero false positives

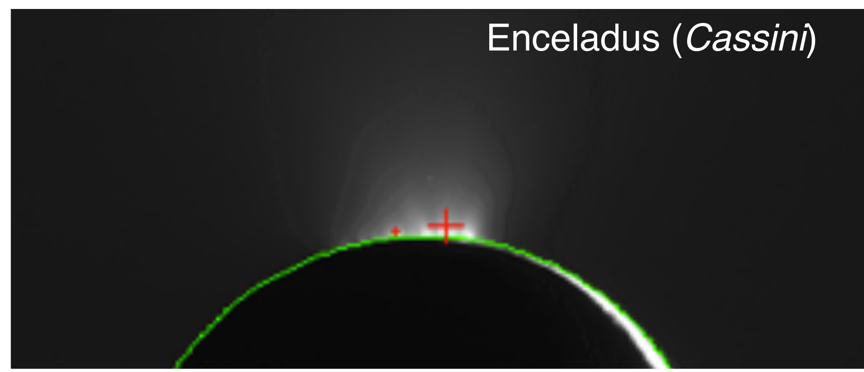

Image annotations from Thompson et al., PSS 2012. Original Hartley 2 images from EPOXI, courtesy NASA. Original Enceladus Image from Cassini mission, courtesy NASA.

A Step further: Sensor Webs

(M.C. Burl, M. Garay, et al)

- Timeliness – respond quickly to short-lived events
- Deficiency – overcome limitations of individual sensing agents
- Provide rich multi-modal observations, particularly of objects that evolve in space and time, such as clouds.
- Object-centric datasets

Geostationary Satellites
GOES-West

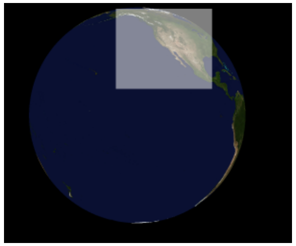

+ wide area coverage
+ dwell over one location - persistence
+ dense temporal sampling (15-30 min)
- lower spatial resolution
- lack of specialty instruments

Polar Orbiters

Terra et al.

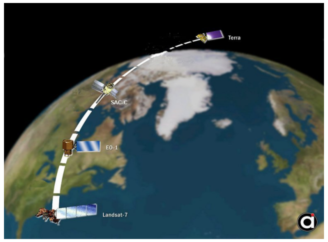

Aqua et al

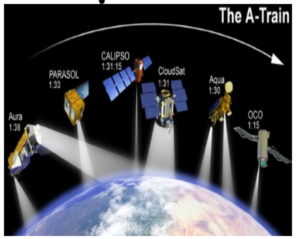

+ high spatial resolution
+ specialty instruments
  Lidar, Multi-angle imager, Cloud Profiling Radar
- cannot dwell over one location
- infrequent revisits (16 day repeat)

Combine into Sensor Webs to Overcome Individual Weaknesses and Exploit Strengths

Adaptive Sky Demonstration Overview

(M.C. Burl, M. Garay, et al)

Bezymianny Volcano
55.978°N,160.587°E,
Kamchatka

Earthquake:
2007/10/14
14:37:05 UTC

Seismic Signal Trigger

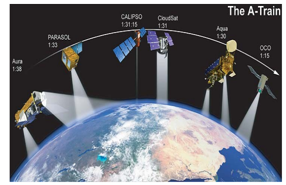

Check for Next EOS Overpass

A-Train Overpass:
15:35 - 15:40 UTC
(Nighttime)

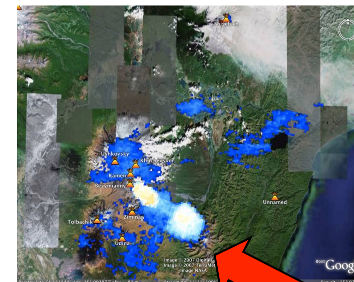

Plume Evident in MODIS Band31-Band32 Signal

Initiate Adaptive Sky Feature Tracking using GOES BTD* Data

2007/10/15

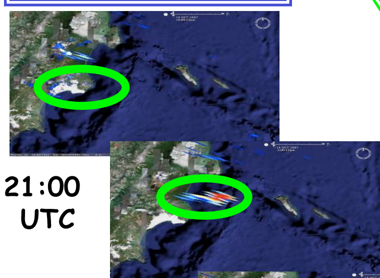

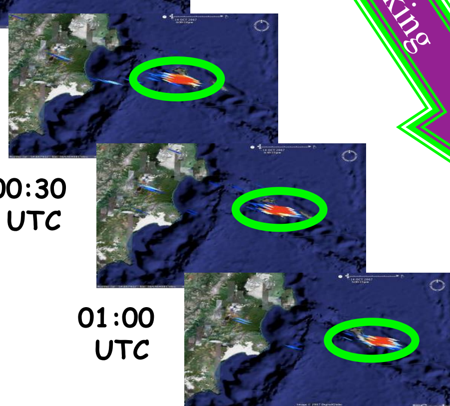

18:00 UTC

21:00 UTC

00:00 UTC

00:30 UTC

01:00 UTC

GOES

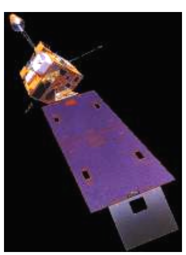

Feature Tracking

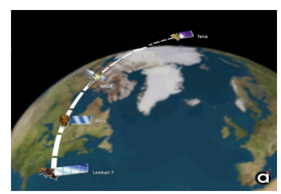

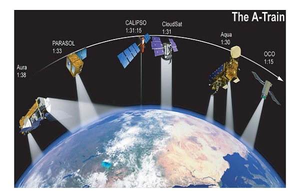

Check for footprint collisions between EOS instruments and tracked ash clouds

Acquire additional observations of ash clouds using the specialized, high-resolution instruments on the EOS polar orbiters

*BTD = Brightness Temperature Difference

Earth Science Technology Office

Follow-up Observations of Ash Cloud 3

(M.C. Burl, M. Garay, et al)

Terra Overpass
2007-10-16T00:05 UTC

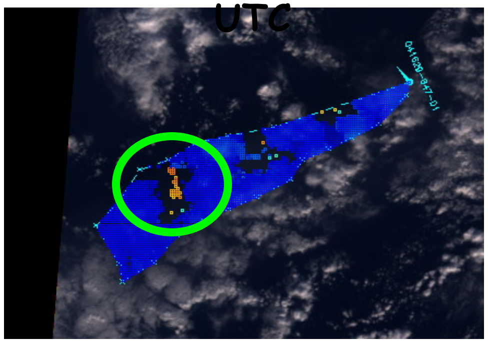

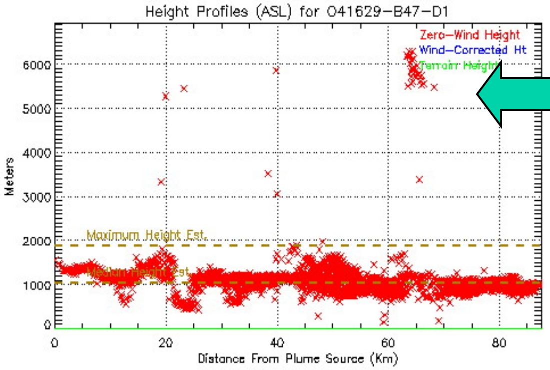

MISR Height Profile

•MISR Stereo Heights indicate a Cloud at ~6 km, with lower clouds at 1-2 km.
•MISR Aerosol Retrievals indicate non-spherical particles in this region, consistent with ash

A-Train Overpass
2007-10-16T01:50

•CALIOP lidar indicates an extremely thin aerosol layer at an altitude of ~6 km in the region.
•The CloudSat radar does not have any returns in this area, indicating extremely small particles.

Conclusion

Adaptive Sky feature tracking allowed observations made in mid-ocean to be associated unambiguously with an ash cloud from the Bezymianny eruption, even with a time difference of ~20 hrs and a spatial separation of ~400 km.

First observations of a volcanic ash cloud from the CALIOP lidar on CALIPSO. Without tracking through the GOES BTD sequence, the returns would have been attributed to cirrus clouds instead.

MISR stereo-derived heights for the ash cloud can be compared directly to the CALIOP lidar heights; MISR aerosol product lends confidence to the assertion this is indeed an ash cloud.

Satellite + CuPIDO Observations

(M.C. Burl, M. Garay, et al)

Overview: The Cumulus Photogrammetric, In-Situ and Doppler Observations (CuPIDO) field program was carried out in summer 2006 near Tucson, Arizona.

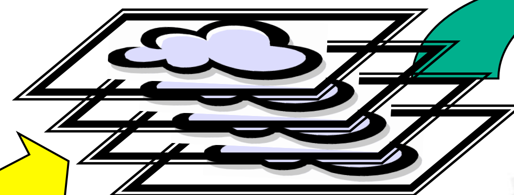

“Top” view from MISR

Reconstructed 3-D* Cloud Volume

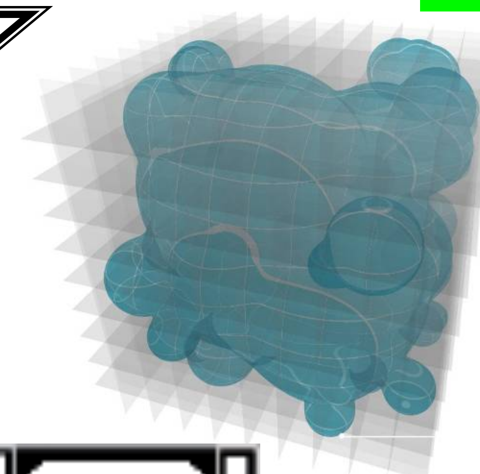

MISR Instrument on NASA EOS-Terra

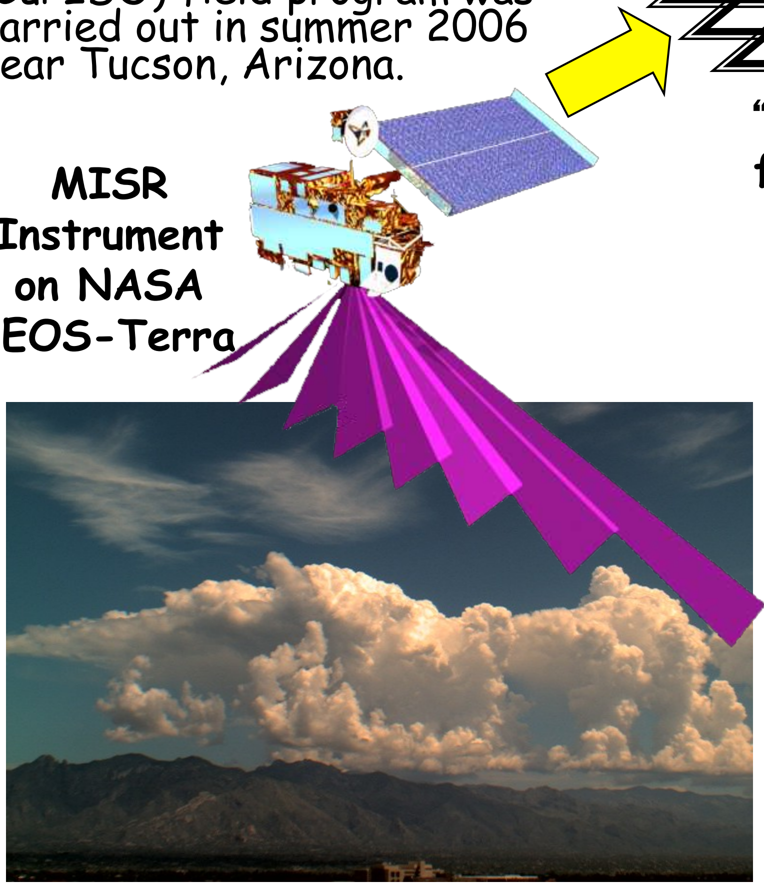

Continuously Sampling Stationary Camera

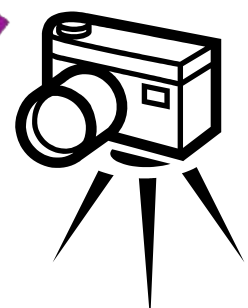

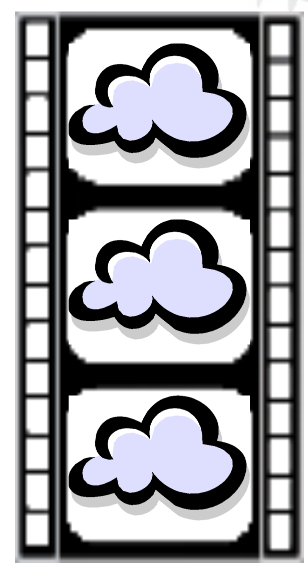

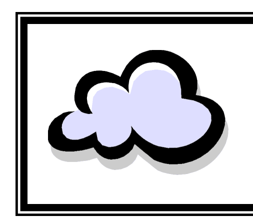

“Side” view from camera

Operational Concept Vision: Satellite + CuPIDO

*3-D = three-dimensional

The Value of Long-term Persistence

Stabilized GOES-West Imagery

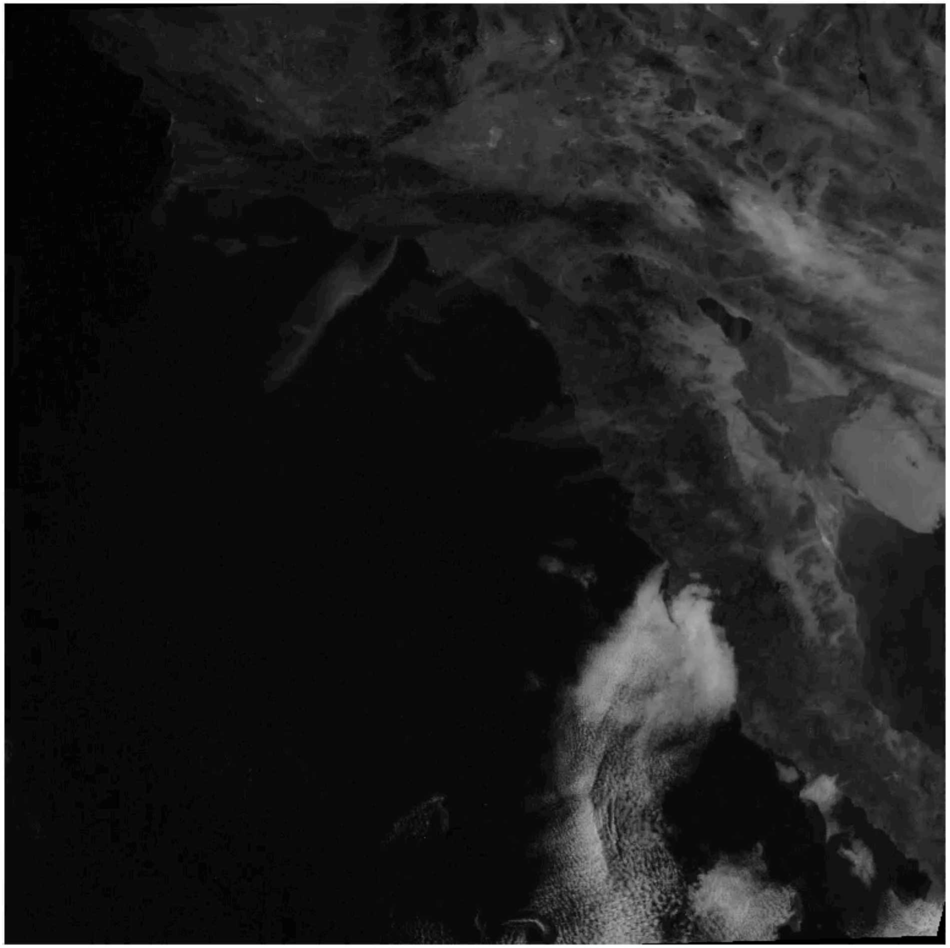

Southern California Wild Fires/Smoke Plumes

(Y. Wang, M.C. Burl, M. Garay)

Summary

Gazing will produce a flood of data!

- Need automated methods for analysis and cataloging feature locations.
- From stabilization/registration to 3D recovery to recognition/discovery

Height Recovery adds new information that can simplify detection/recognition

- Example using MOLA data for crater detection
- Using multi-angle height recovery to improve car detection is promising but difficult (at the limit)

Discovery

- Best approach under certain conditions – unknown or highly variable appearance
- Height and temporal change detection from staring add valuable dimensions for discovery

Triggering/Event-based Observations

- Pointing when and where something interesting is happening/about to happen

Sensor Webs

- Combine observations from multiple sensors to overcome weaknesses of each.
- Collective capability is greater than sum of parts.
- Persistence extremely valuable for some types of observations.