Jet Propulsion Laboratory
California Institute of Technology

"Look and you will find it – what is unsought will go undetected." - Sophocles

COSMIC

Content-based On-board Summarization to Monitor Infrequent Change

PI: Dr. Lukas Mandrake
Co-I: Gary Doran (398), Masha Liukis (398), Steven Lu (398),
Umaa Rebbapragada (398), Jimmie Young (398), Kiri Wagstaff (398)
Interns: Aaron Alvarado, Valentin Bickel, Anneliese Braunegg, Tina Cao, Paul Horton, Daniel Jeong, Michael Munje, Asher Trockman

RPC 2020 - Virtual Research Presentation Conference

Program: Strategic Initiative (6x)

Assigned Presentation #18

Executive Summary – Global Change Monitoring

- Spacecraft like the Mars Reconnaissance Orbiter (MRO) scan the entire surface of Mars every few days at high resolution
- Yet only a handful of pre-selected images per day can be sent back to Earth due to limited communications bandwidth
- Transient science (active surface change) is rarely captured… can't know where/when it will be to request image
    - Fresh Impacts, rockfall, dust devils, araneiforms, active gulleys, recurring slope linea, active transverse aeolian ridges, dust uplift events, etc.
- Conventional change detection requires a before and after picture for differencing
- Global case would need the whole planet stored onboard as pixels for comparison… would take XX TB of onboard storage!
- **Innovation**: Use (2) onboard science autonomy capabilities to globally capture & monitor high-impact transient science targets

1) Store highly summarized representations of surface (rather than raw pixels) onboard and identify change in summary
    - Agnostic to expectation, can find new kinds of change at many spatial scales
2) Directly recognize known kinds of active/recent transient events in still images
    - Identify, map, and return cutouts of hi-res data only around targets of interest

- **Outcome**: Demonstrated utility of system via "Mission into the data" and discovered ~100 new, previously unknown fresh impacts in the existing MRO CTX image collection

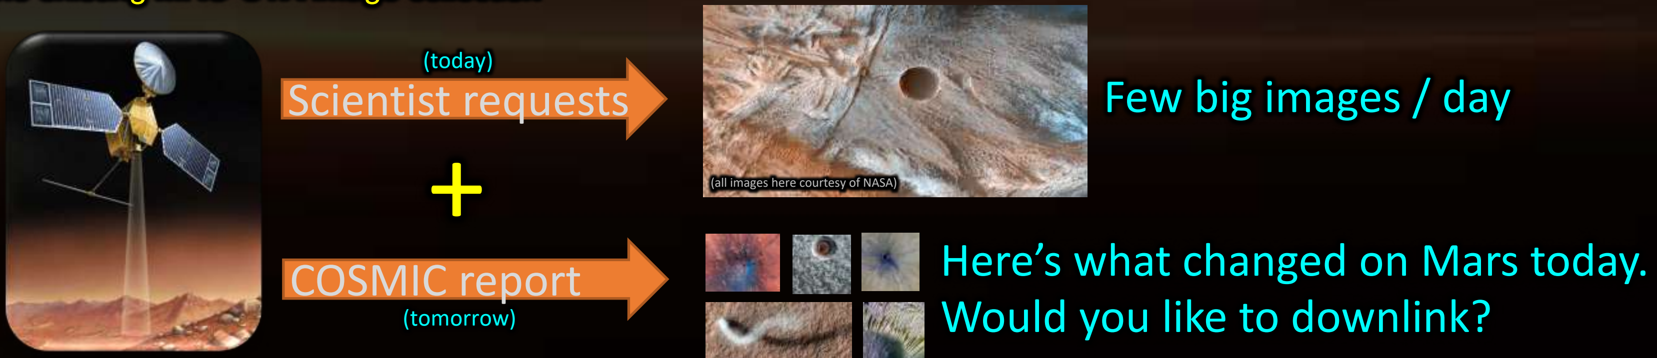

Challenge: Changing Planet, minimal downlink

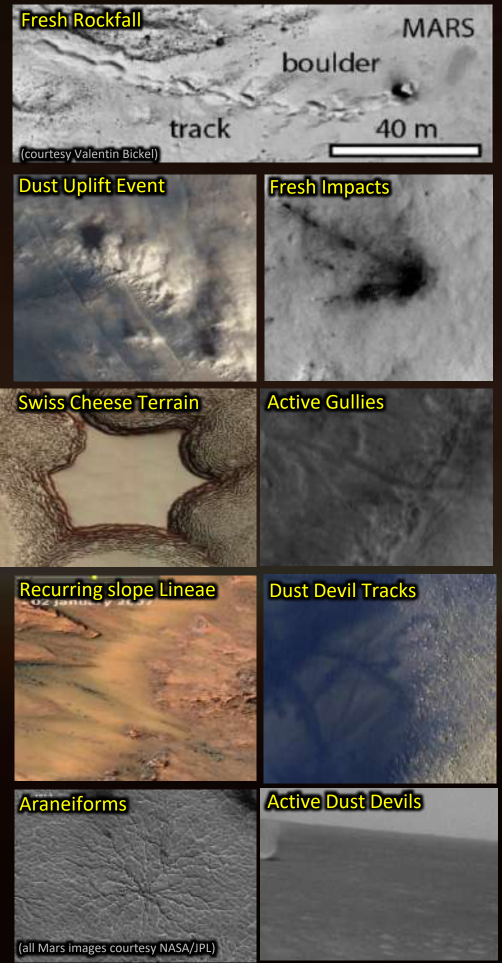

Why transient science matters?

- Seismic activity
- Regolith geomechanical properties
- Active aeolian processes / wind characterization
- Replenishment / transport of atmospheric dust
- Subsurface composition
- Modern planetary bombardment rate
- CO_2 & water ice transport via sublimination and melting
- Dry vs. wet erosive processes
- Broader implications for planetary chronology

Direct change detection on Earth? Simple!

- Align two images pixel-for-pixel
- Difference them
- Look for change… voila!

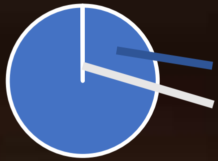

For a single 17-day MRO repeat cycle

Single global mosaic of Mars ~400 TB (~2 Gbps)

Downlink possible ~0.3 TB (~1.5 Mbps)

State of Art: Exhausting Manual Search

Missed findings, unscalable human effort

Planetary Change Detection? Hard.

- Can't send back whole planet's images
- Especially repeatedly via overpasses
- Can't store whole planet onboard
- Hence, direct change detection won't work
- How to recognize, capture, and characterize global Martian change?

Innovation: Onboard Summarization for Science

(courtesy Pixabay free graphics)

(courtesy NASA)

(courtesy Pixabay free graphics)

(courtesy NASA)

(courtesy COSMIC project)

(courtesy NASA)

(courtesy NASA)

(courtesy PixaBay free graphics)

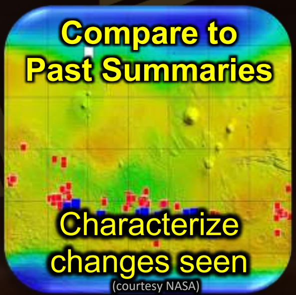

(courtesy NASA)

Innovative Concepts

- Leverage Upcoming Space Computing
- Leave camera on & process everything
- Trustable science autonomy
- Store content-based summarizations
- Perform change detection between summaries
- Can now trade onboard compute for bandwidth!
- Many kinds of change can be captured
- Camera + compute + COSMIC is new smart instrument!

Scientists in Control

- Scientists reconfigure COSMIC on-the-fly
- COSMIC sends reports on global change
- Scientists more informed to react:
  - Request targeting of observed change?
  - Request tiny snapshots already waiting?
  - Request high-res monitoring of a site?
- Miss a comm pass? COSMIC keeps working.

Change Detection Capabilities & Progress

Known Science Targets
(surface classification)

- Know what, but not where/when
- Can look for any surface feature
- Needs example images for training
- Alerts provide time/space map
- Works on entire globe

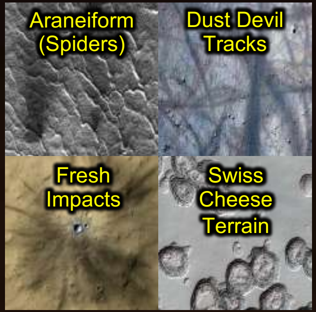

TRL ~5

- Tested on entire CTX database
- Actual science discoveries attained!
- Related task infused into PDS archive
- Ready for flight-oriented maturation

Agnostic Change Detection
(image-summary-based change)

- Don’t know what, where, or when
- Discover new transients
- Global reports on change characterization
- Agnostic to expectation or known targets

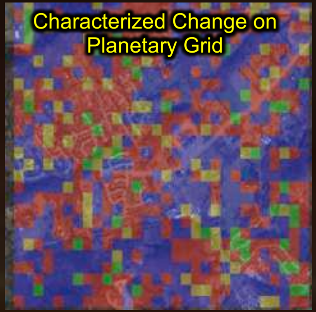

TRL ~3

- Feasibility study complete on real data
- Computational goals met
- New change-detection req: Stability
- Refinement needed for specific science use-cases and rigorous V&V

Region Monitoring
(Pixel change for small areas)

- Know where, but not what/when
- Classical change detection
- Only a ~1k small regions at a time
- Alerts when something starts to happen

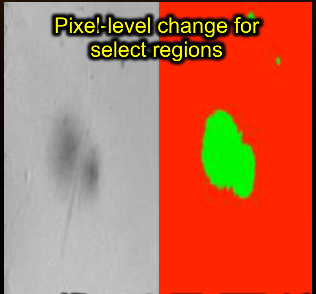

(all images courtesy NASA or COSMIC project)

TRL ~3.5

- Excellent performance on small tests using real & synthetic data
- Tolerant to expected operational environment
- Ready for rigorous V&V

Results

Protecting the Autonomy: Initial Image Filtration

Goal: Filter inputs that are either uninformative or problematic for downstream algorithms with a sequence of classifiers built for each problem class

Requirements: Low false positive rate in order to minimize missed science

Examples of Problematic Input from HiRISE

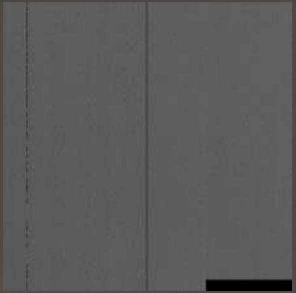

Dark Images / Artifact

Dark images contain no information. Artifacts will trigger false positives from downstream algorithms

Dust

Atmospheric dust obscures Martian surface

Miscalibration

Vertical striping caused by miscalibration will trigger false positives from the downstream algorithms

Featureless

Uniformly featureless landscapes are uninformative and may trigger false positives from downstream algorithms

Evaluation: 260 image HiRISE test set with an equal number of classes. Inputs ran sequentially through the system. Positive classifications rejected and negatives allowed through.

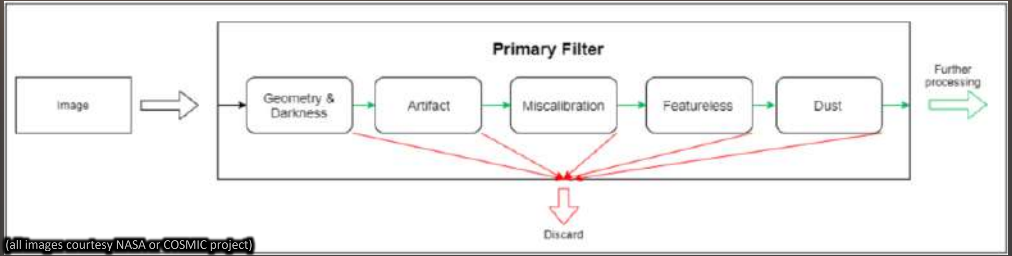

(all images courtesy NASA or COSMIC project)

Results

Accuracy:	0.91
False Negative Rate:	0.02
False Positive Rate:	0.10

Localization – Where Are You Looking?

- Images must be **localized** accurately to be compared across multiple observations.
- Spacecraft pointing/position awareness isn't good enough. Prevents any change detection match-up.
- Use content-based, multi-tiered approach to reach accuracy required by each science use case (1-100 m):

Araneiforms (“Spiders”)

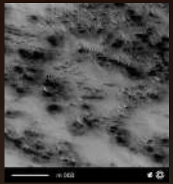

“Swiss Cheese” Terrain

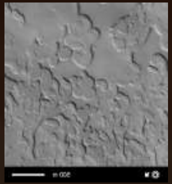

Fresh Impacts

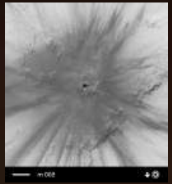

Dust Devil Tracks

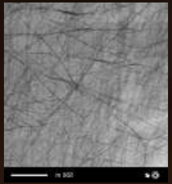

Coarse alignment using tiepoints extracted from CTX observations

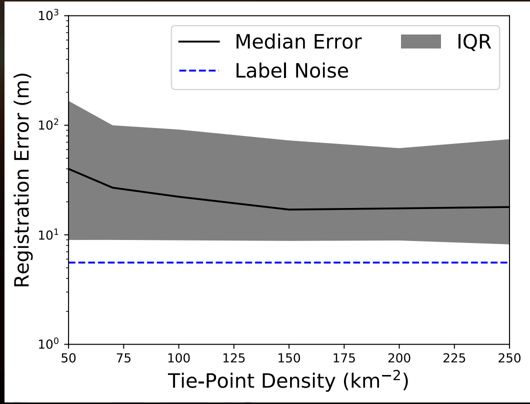

(all images courtesy NASA or COSMIC project)

Novel method for content-based localization

Initial Localization

Use Spacecraft Ephemeris

Error: O(1 km)

Coarse Alignment

Register to Global CTX Mosaic Tie-points

Error: O(10-100 m)

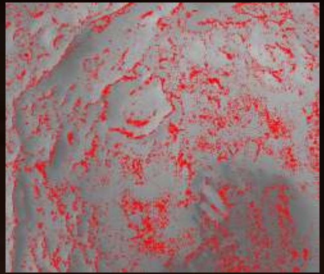

Fine Alignment

Direct Pixelwise Registration/Comparison

Error: O(1-10 m)

Narration: Dr. Gary Doran

Target Monitor: Pixel-based change in select regions

Target Monitoring: when you know *where* to look but not *when*

- Store raw pixels of select regions onboard, perform pixel-level registration to align
- Use a classifier (TextureCam) that uses before/after image pair to detect change
- Pixel-wise classifier uses a context window to tolerate some pixel misalignment
- Send alerts to ground when, after months of calm, something starts to happen

Above: seasonal frost

Right: avalanche in polar region

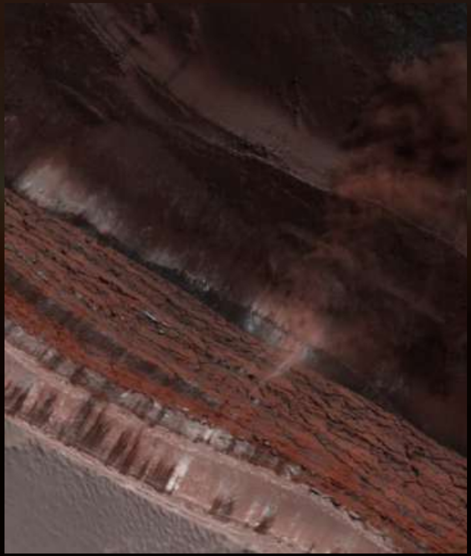

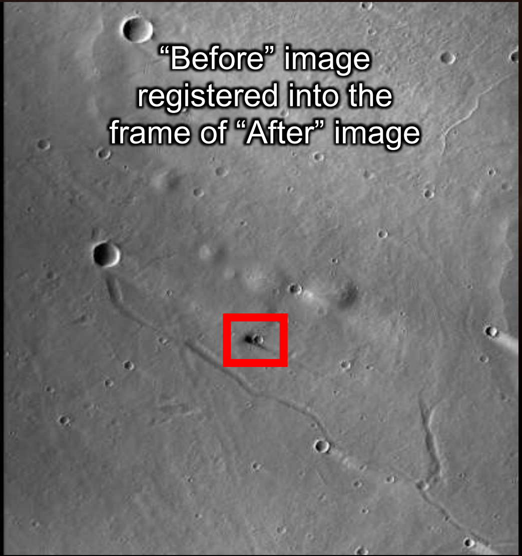

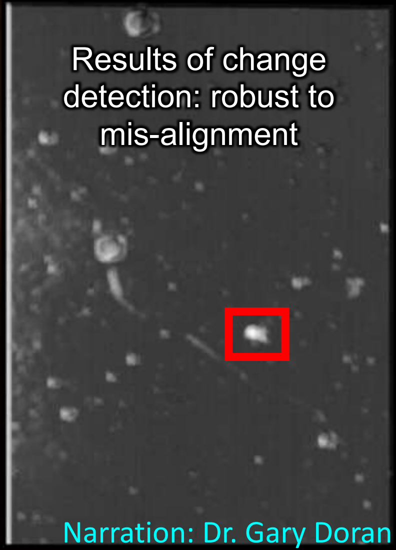

(all images courtesy COSMIC project)

Narration: Dr. Gary Doran

Transient-suggestive Surface Features: Landmarks

"Landmarks" are surface features of interest, including craters, slope streaks, dust devil tracks, etc.

- Several new labeled data sets created with landmarks in CTX (low resolution) and HiRISE (high resolution)
- Enabled assessment of landmark detection methods previously developed (Wagstaff et al., 2012) as well as creation of new landmark classifiers using state-of-the-art technology
- Finding: Random forest achieves same or better performance as convolutional neural network (CNN) on high-resolution HiRISE image tiles, at 50% of the CNN runtime cost. Good news for potential onboard deployment.

Fresh impact found in HiRISE image by Random Forest classifier

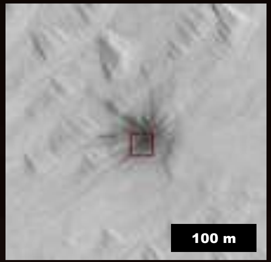

Case study: Deployed "fresh impact" CNN classifier across 112,000 CTX images of Mars (Munje et al., 2020). Currently reviewing candidates in order of probability of fresh impact. Of top 1000, 69% are fresh impacts and 9% were previously unknown (new discoveries!)

- CNN out-performed random forest on CTX image classification, and ground-based analysis can accommodate more costly models
- New discoveries are being submitted for HiRISE close-up imaging (via HiWish)
- Significance: Machine learning helps accelerate scientific discoveries (at current rate: increasing fresh impact catalog by 50%)

Sample from 1M classifier candidates

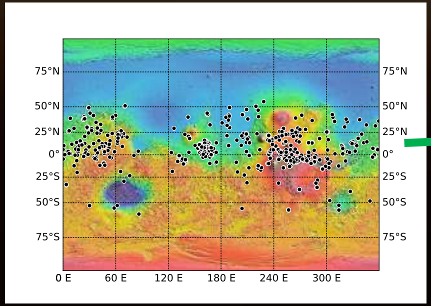

Top 1000 after human review

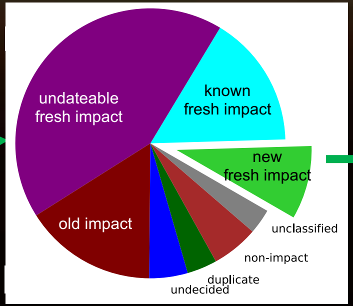

First new discovery found (using CTX)

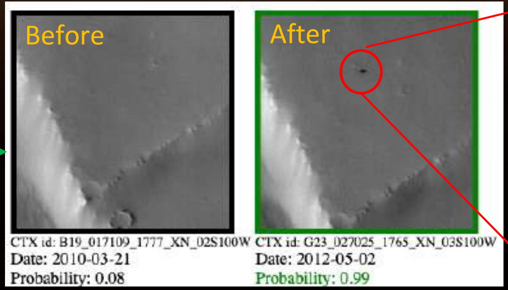

Follow-up HiRISE view

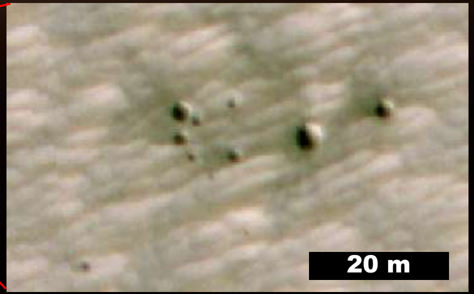

(all images courtesy NASA or COSMIC project)

Narration: Dr. Kiri Wagstaff

Global Agnostic Change Detection

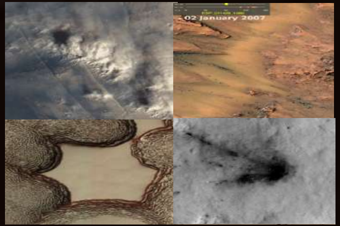

Global agnostic change detection scours the entire Martian surface for any change, without targeting.

- Salience-based approach: Save "interesting" region summaries for later comparison
- Tiling-based approach: Save summaries for whole planetary grid (HEALPix) for later comparison

Salience-based change detection

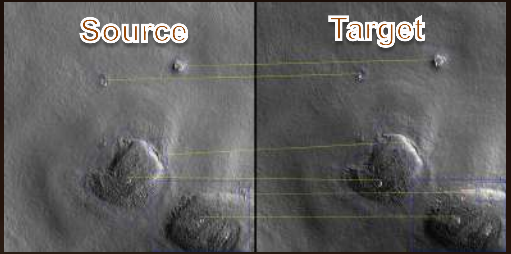

Tiling-based change detection

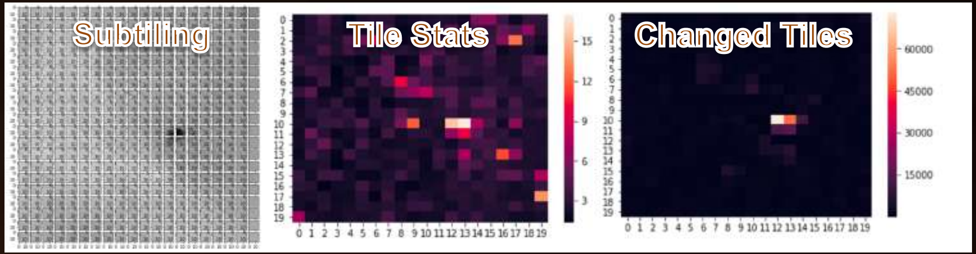

Stability analysis to assess how imaging parameters affect change detection performance; genetic algorithm (GA) to optimize stability.

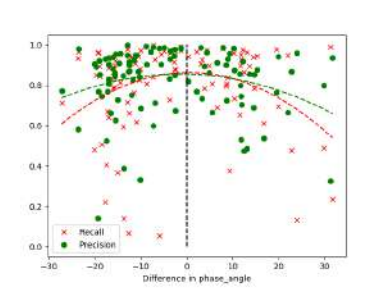

Illumination v.s. performance

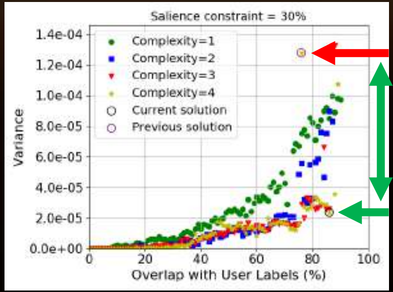

Previous solution

Improvements

Re-optimized Solution

GA Optimization

Narration: Steven Lu

Conclusion: Both salience-based and tiling-based change detection methods are feasible for the onboard settings with improvements to feature representations and stability of salient landmarks detection.

	Storage	Compute	Stability
Salience-based change detection	Only features of salient landmarks will be stored (smaller database)	More compute resources to find salient landmarks	ID of salient regions over-sensitive to unimportant changes
Tiling-based change detection	Features of all sub-tiles will be stored (larger database)	Tiling algorithm is computational efficient	Robust to small errors in localization

Custom Science Payload Example: FOLD

COSMIC framework can accommodate custom science algorithms.

Example: FOLD (Fast Oriented Layer Detector)

Originally designed for analyzing rover images to find layered rock, the same algorithm can identify layers within orbital imagery.

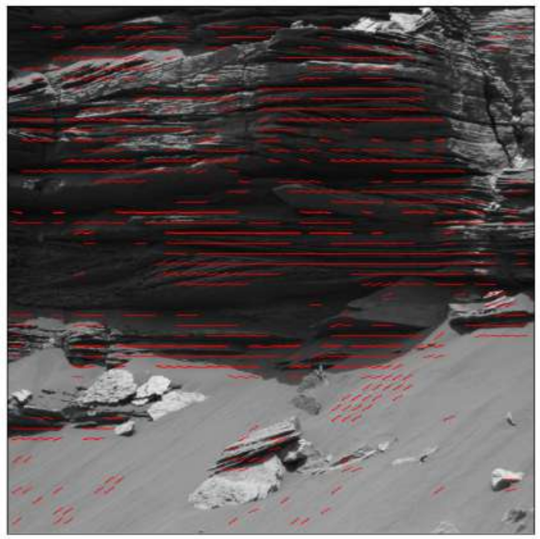

(all images courtesy NASA or COSMIC project)

Above: layers detected in a Mastcam Image

Right: layers detected from orbital CTX imagery

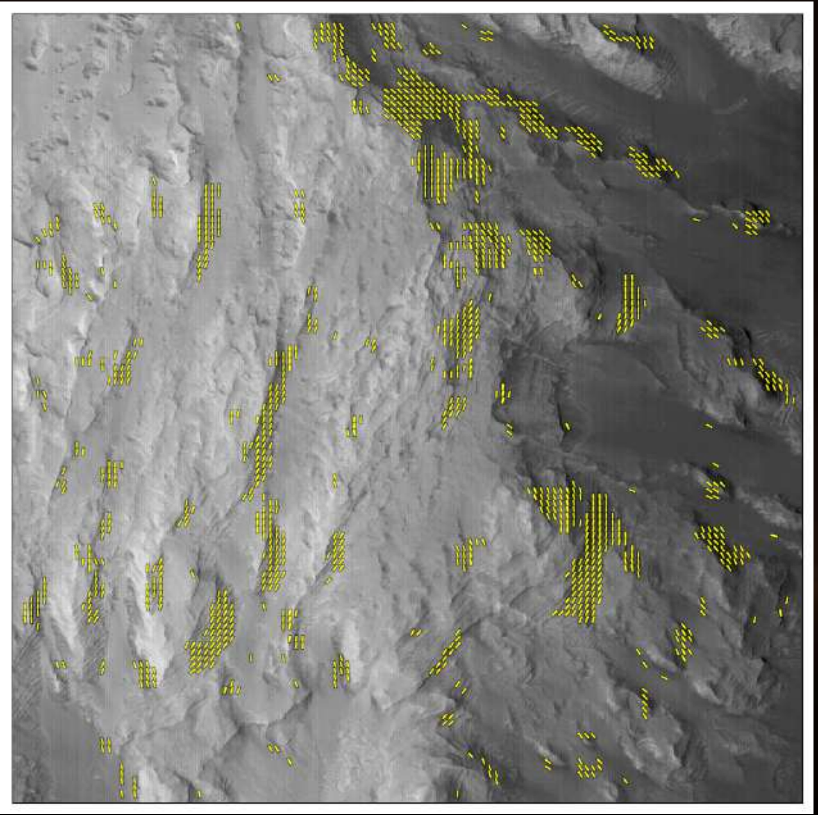

Narration: Dr. Gary Doran

On-Demand Global Mars Database

Onboard Databases of Current Martian Information

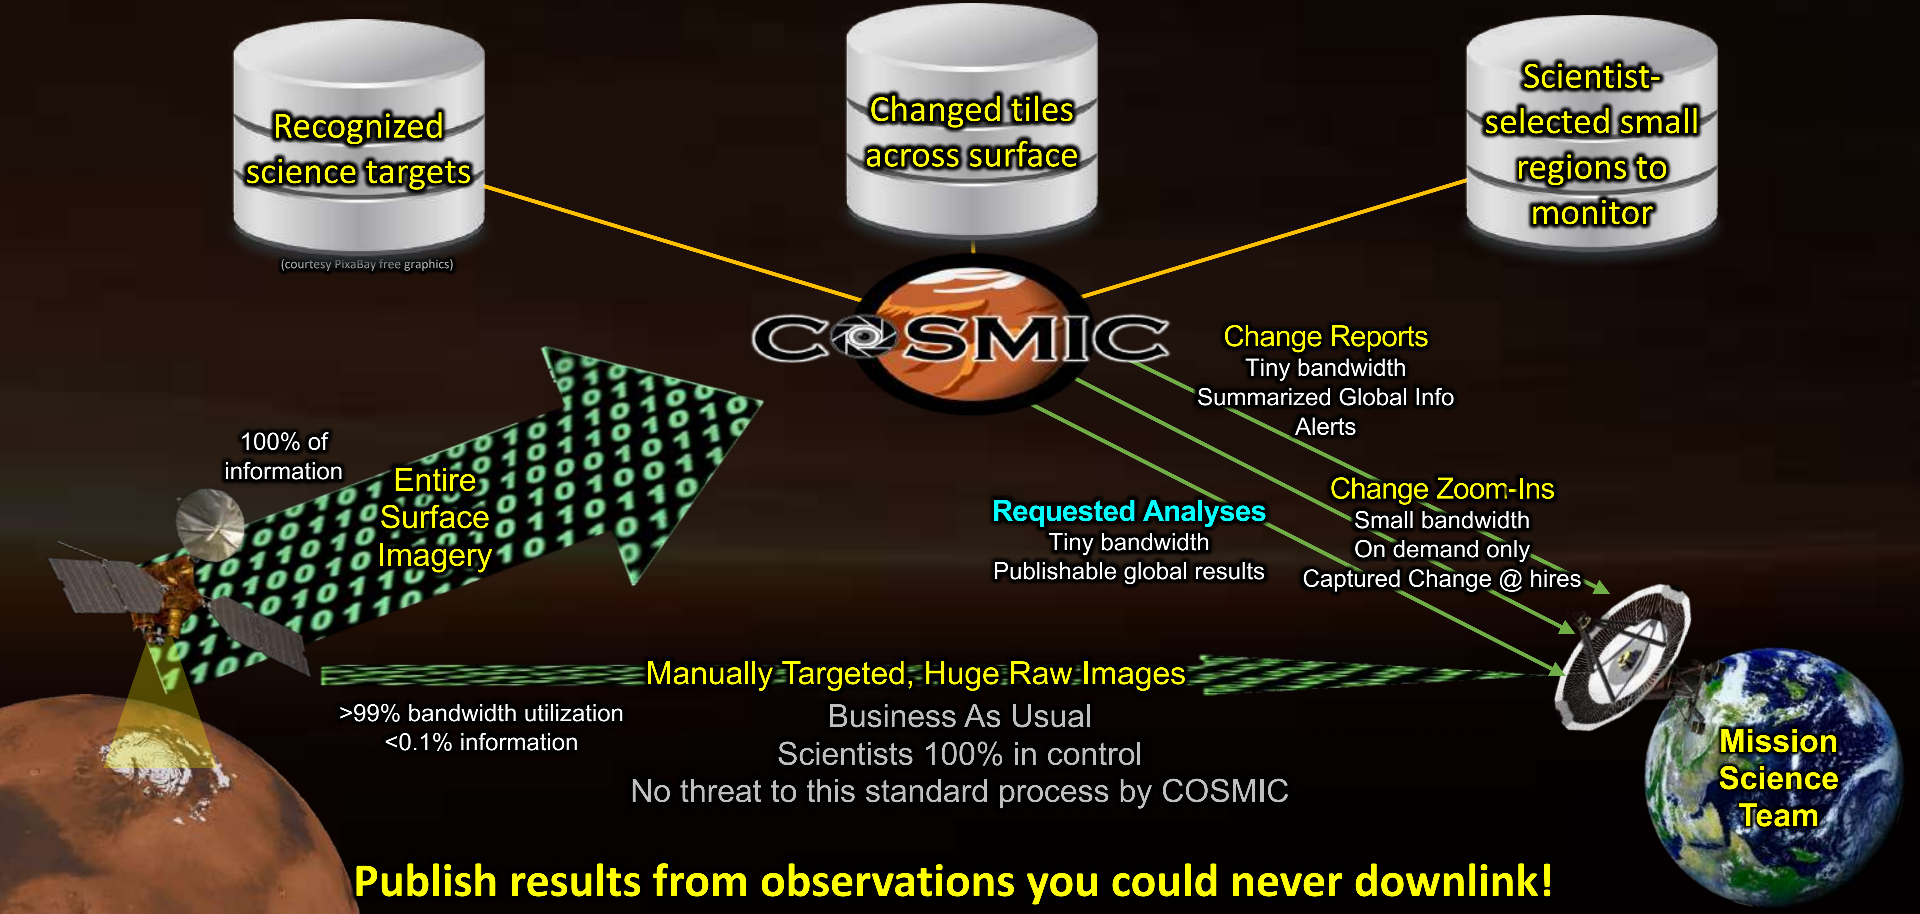

Final Status & Future Plans

Ambitious original goal: all functionality + flight-like system implementation + extensive validation (aim high!)

Successes:

- Identified / retired significant risks for onboard change detection
    - Stability (underlying systems need more than accuracy for change; they must also say the same answer, right or wrong, every overflight)
    - Targeted surface feature detection (Demonstrated **real science yield**, accuracy levels, and stability appropriate to deployment)
    - Tile-based change detection (Feasibility confirmed, requires additional research to summarize tile contents efficiently)
    - Salient-region change detection (Identified as not yet sufficiently stable for change detection purposes)
    - Localization (identified as major dependency, addressed by new content-based, achieved required performance)
    - Onboard storage (identified a few TB onboard storage sufficient to enable global change detection)
- Demonstrated ACTUAL science yield during SRTD
- Contributed heavily to KISS Nebula workshops (superset of COSMIC capabilities)
- Ongoing FY21 formulation plans
    - (formulation) Combine COSMIC + sister SRTD tasks + KISS Nebula workshop into JNEXT submission: a global change monitoring smart camera
    - (pending) Submitted NASA Game Changing Development proposal
    - (pending) NASA AMMOS tile-based change detection system
    - (not selected) Submitted Topical RTD for Onboard Cartographer localization

Publications & Press Releases

Published Papers & Presentations

Gary Doran, Steven Lu, Maria Liukis, Lukas Mandrake, Umaa Rebbapragada, Kiri L. Wagstaff, Jimmie Young, Erik Langert, Anneliese Braunegg, Paul Horton, Daniel Jeong, and Asher Trockman. COSMIC: Content-based Onboard Summarization to Monitor Infrequent Change. Proceedings of the IEEE Aerospace Conference, 2020.

Michael J. Munje, Ingrid J Daubar, Gary Doran, Kiri Wagstaff, and Lukas Mandrake. Large-Scale Automated Detection of Fresh Impacts on Mars Using Machine Learning with CTX Observations. 11th Planetary Crater Consortium, Abstract #2065, August 2020.

(several datasets and additional manuscripts in process)

Significant NASA/JPL press release resulting in dozens of featured articles:

- "Machine Learning Software is Now Doing the Exhausting Task of Counting Craters On Mars," Universe Today
- https://www.universetoday.com/148165/machine-learning-software-is-now-doing-the-exhausting-task-of-counting-craters-on-mars/

- "NASA Unveils a Cluster of Mars Craters Discovered by AI" on Analytics Insight
- https://www.analyticsinsight.net/nasa-unveils-cluster-mars-craters-discovered-ai/

- Interview by Austin Cross from KPCC (NPR)

- "Scientists are Finding a Ton of New and Mysterious Craters on Mars" on vice.com
- https://www.vice.com/en/article/4ayjkq/scientists-are-finding-a-ton-of-new-and-mysterious-craters-on-mars

- "NASA uses AI to find new craters on Mars" on Clarksville Online
- https://www.clarksvilleonline.com/tag/hirise/

- Follow-up image on HiRISE database:
- https://www.uahirise.org/ESP_066018_1765

End

Virtual Research Presentation Conference

Content-based Onboard Summarization to Monitor Infrequent Change (COSMIC)

Principal Investigator: Lukas Mandrake (398)
Co-Is: Gary Doran (398), Masha Liukis (398), Steven Lu (398), Umaa Rebbapragada (398), Jimmie Young (398), Kiri Wagstaff (398)
Program: Strategic Initiative (6x)
Assigned Presentation #

Jet Propulsion Laboratory
California Institute of Technology