Risk-aware Autonomy for Resilient Space Exploration

Airbus Meeting

Hiro Ono
Jet Propulsion Laboratory, California Institute of Technology
ono@jpl.nasa.gov

9/13/16

Why Risk-Ware Autonomy in Space?

- Current ground-in-the-loop operation too slow
  - E.g., Mars rovers: <100 m/Sol (<50 m/Sol most of the times)

- Most accessible locations visited
  - Next frontier: subsurface

Caves on Moon and Mars

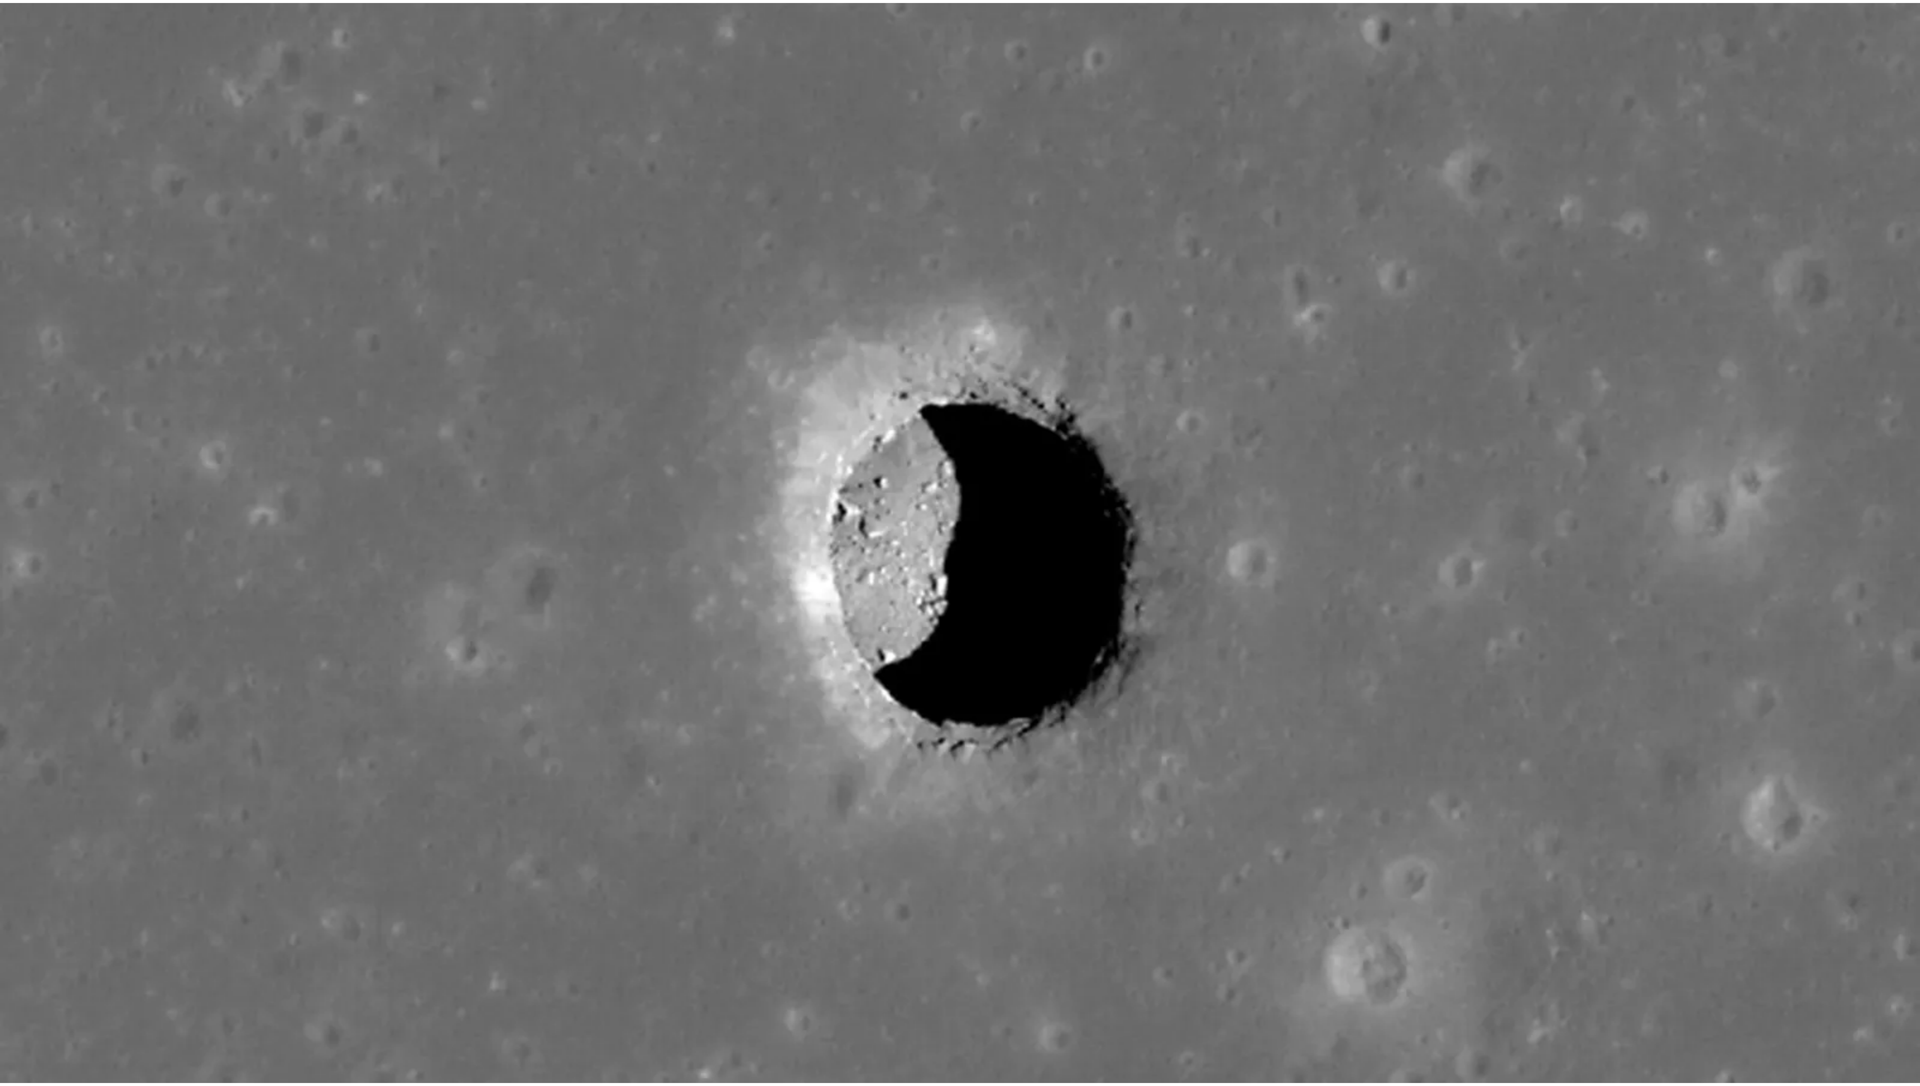

Caves on Moon and Mars

In 2005, when *Cassini* visited a small Saturnian moon, Enceladus…

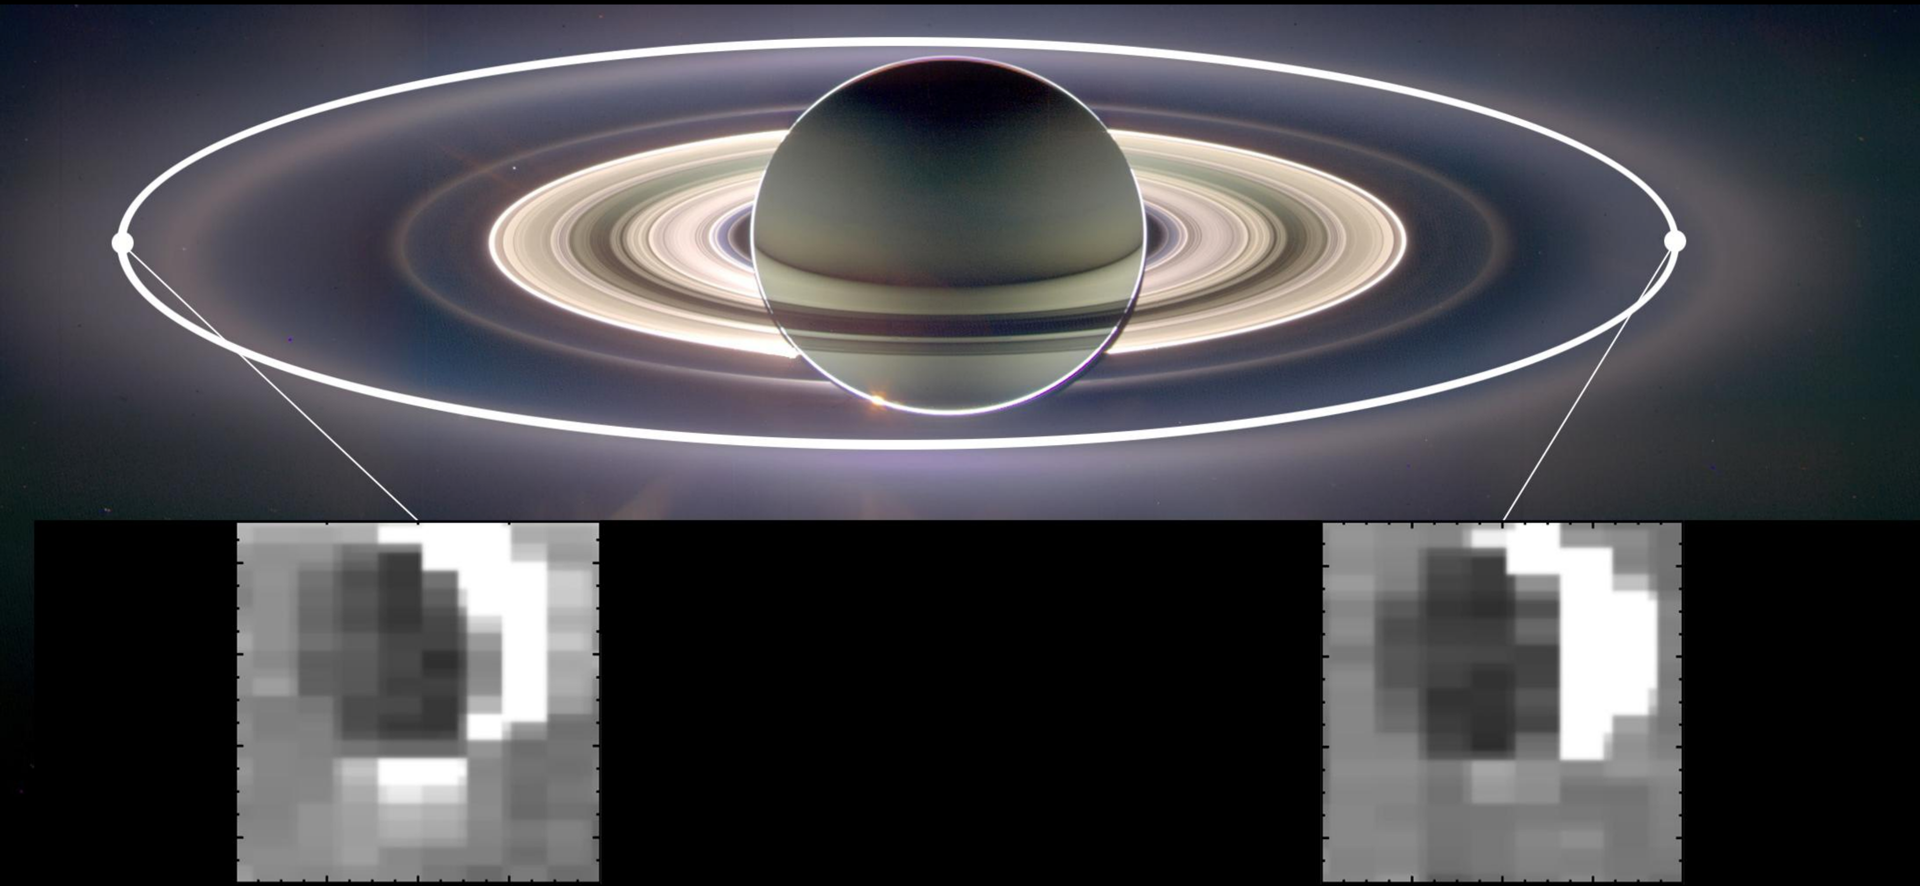

Subsurface Ocean of Icy Moons

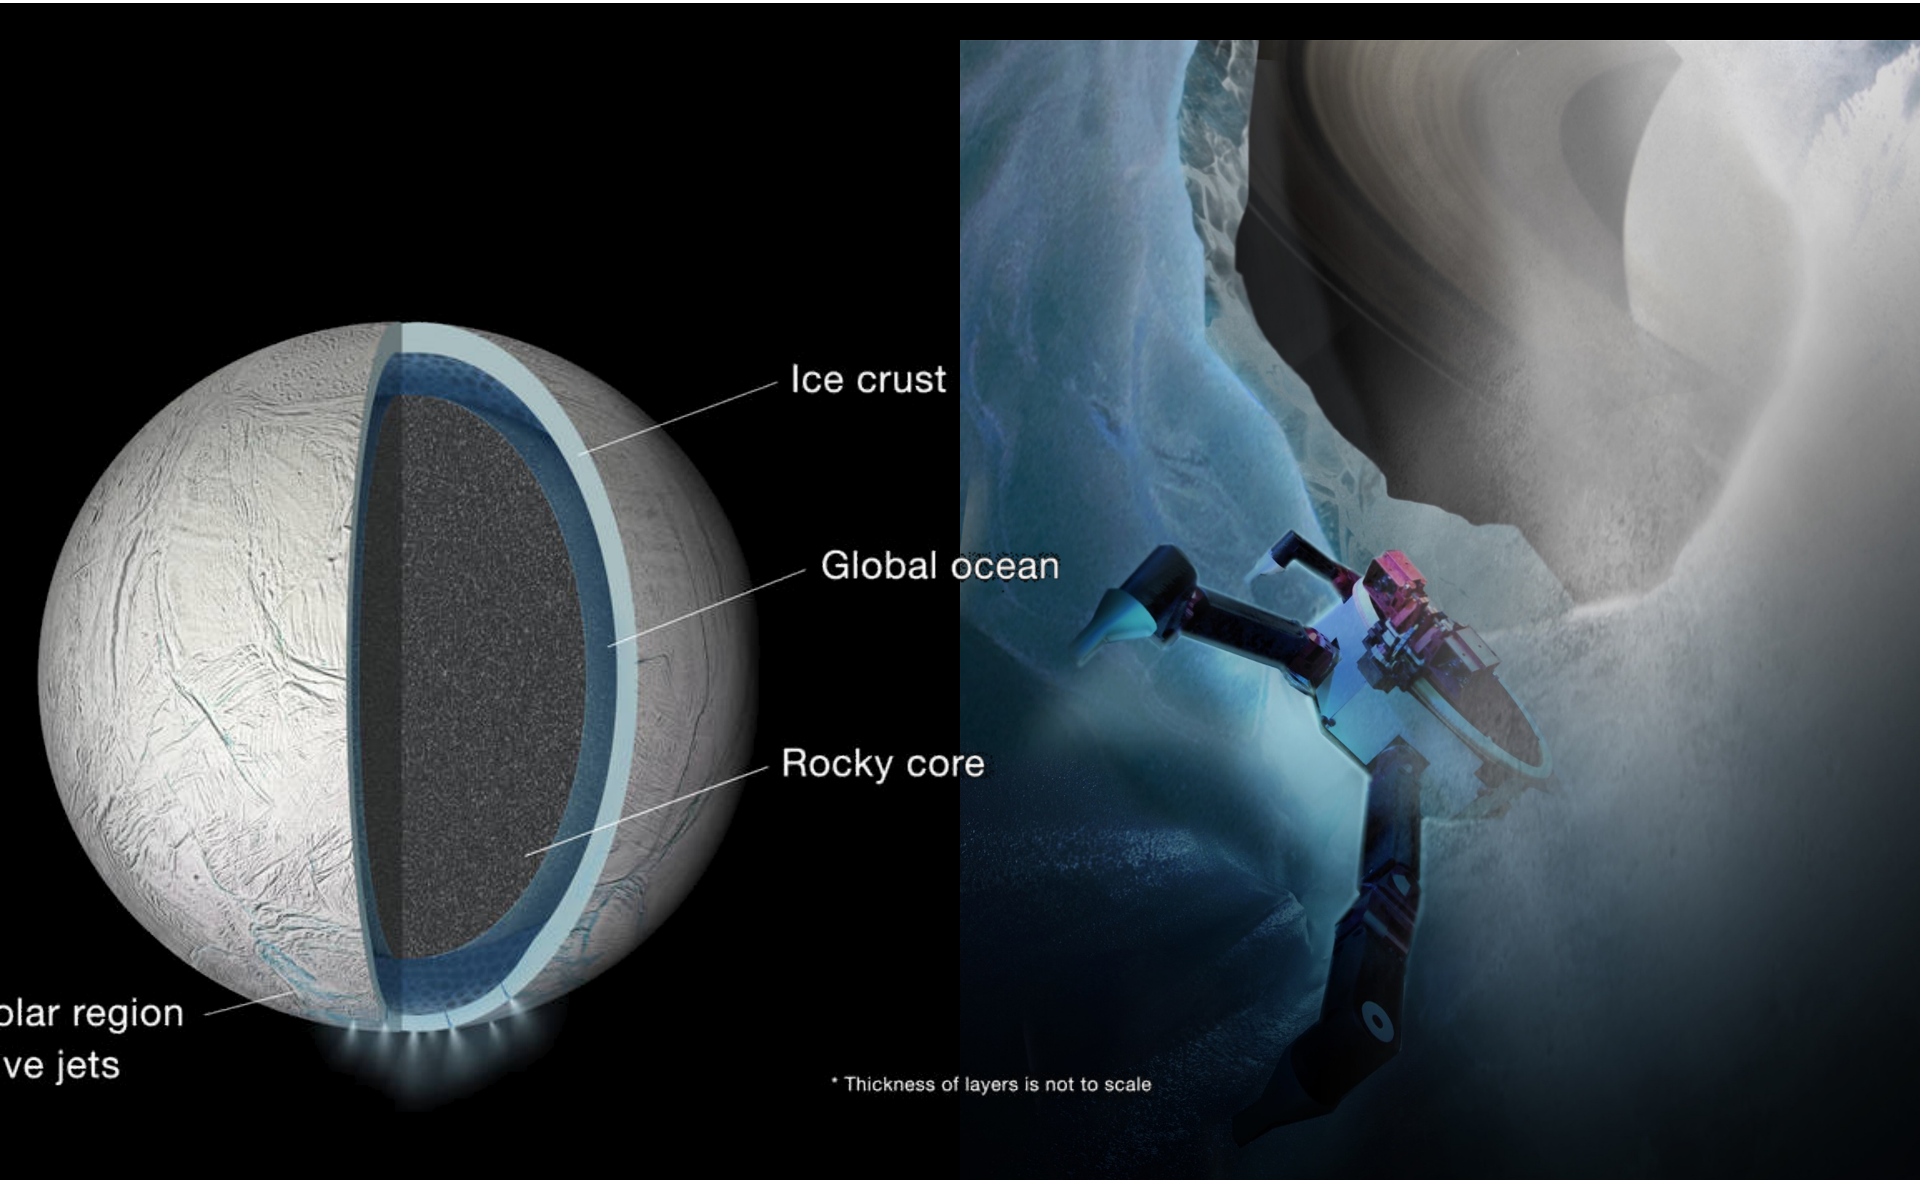

Challenges for Subsurface Exploration

- No orbital reconnaissance
- Limited communication (radio does not propagate in water)
- Limited visibility
- Risk, uncertainty, unknowns

Risk-Aware Autonomy: Overview

Autonomy adapts its behavior depending on the acceptable level of risk

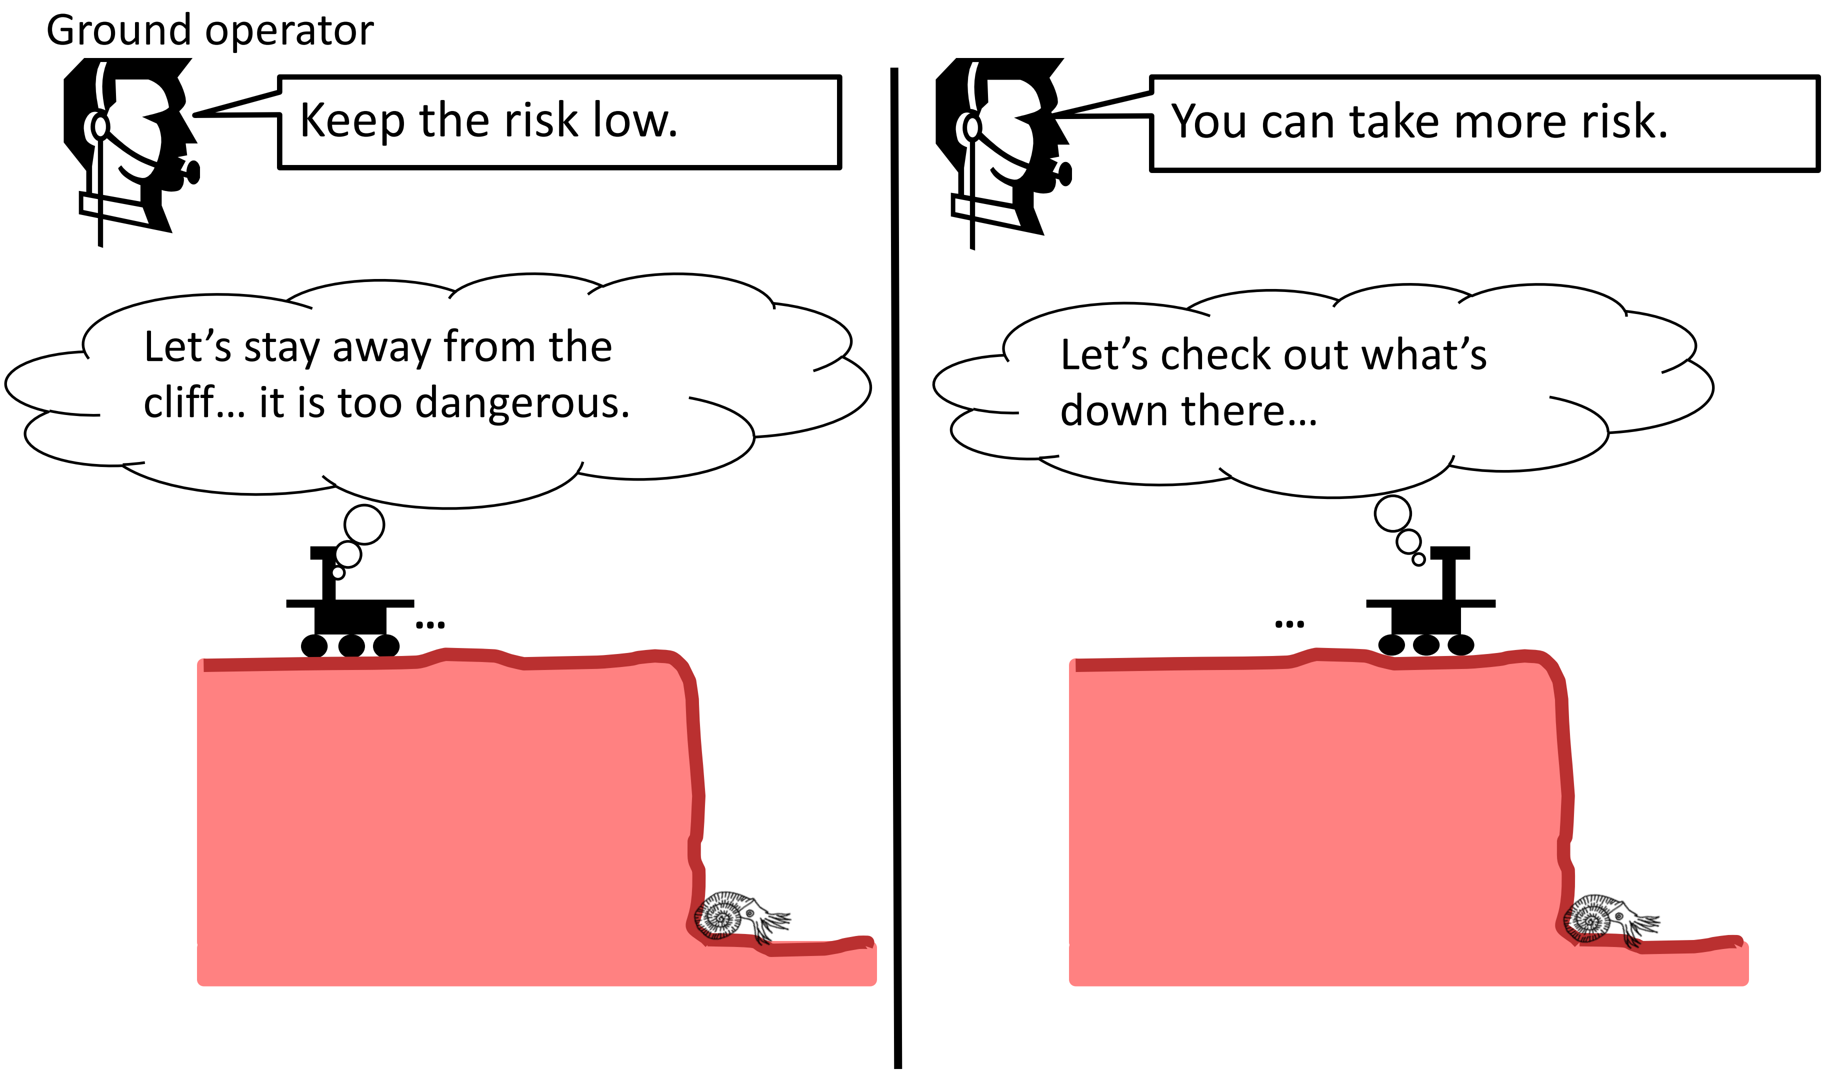

Demo

Low risk threshold

High risk threshold

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Resilient Risk-Aware Autonomy for the Exploration of Uncertain and Extreme Environments*

FY15-16 RTD/KISS

Co-PIs: Mitch Ingham and Richard Murray

Problem being solved:

Develop *Resilient Spacecraft Executive* that:

- adapts to component failures to allow graceful degradation
- accommodates environments, science observations, and spacecraft capabilities that are not fully known in advance
- makes risk-aware decisions without waiting for slow ground-based reactions

Why this is important to JPL:

- Enables robotic explorations of harsh, remote, and inaccessible destinations, e.g., Venus and KBOs
- Reduces operation cost without reducing safety

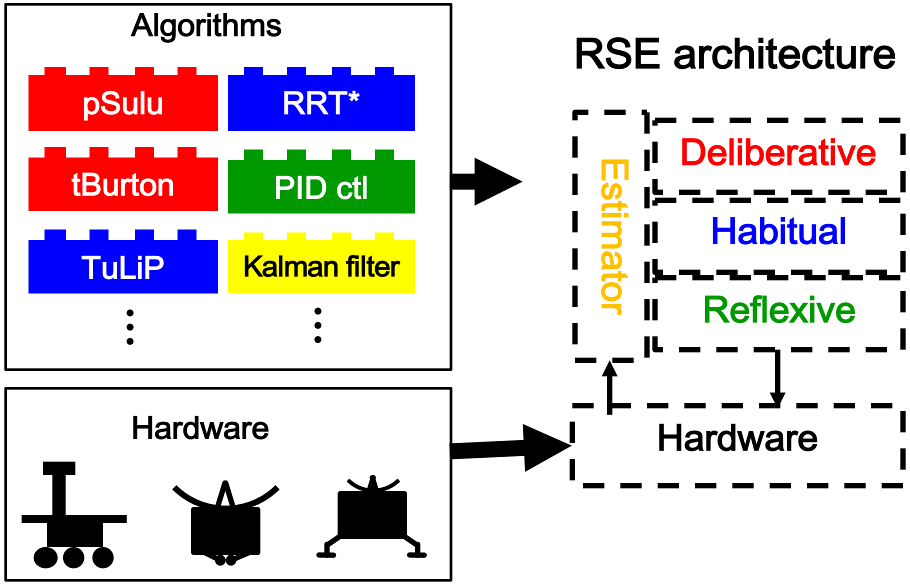

Demo on ATRV rover and ROAMS simulator

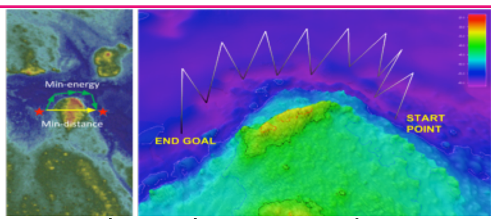

AUV demo by MIT and WHOI

Team Members

Role	Name	Sec.
PI	Mitch Ingham	313
Co-I	Tara Estlin	397
Co-I	Masahiro Ono	347
Co-I	Leslie Tamppari	322

External collaborators

Prof. Richard Murray
(Funded by KISS)

Prof. Brian Williams
(Funded by KISS)

Dr. Richard Camilli
(Funded by KISS)

FY15 Accomplishments:

- •Identified the science drivers for RSE capabilities to be demonstrated for a selected set of reference missions;
- •Developed an architecture specification for the RSE and modeled the architecture formally in SysML
- •Developed and integrated the initial algorithms used in the architectural layers of RSE;
- •Created an integrated RSE capability ready for deployment in software simulation and on rover/AUV hardware platforms that is compatible with the Robot Operating System (ROS) middleware software;
- •Deployed RSE on multiple rovers and AUV systems, in simulation and on hardware.

Risk-Aware, Human-Cooperative Autonomy

Sponsor: ONR PI: Masahiro Ono (347E) ono@jpl.nasa.gov

Project objective: Develop an autonomy that:

- Makes optimal decisions based on risk
- Manages risks at multiple levels of abstraction
- Shares the responsibility of managing risks
- Combines the strengths of each other synergistically

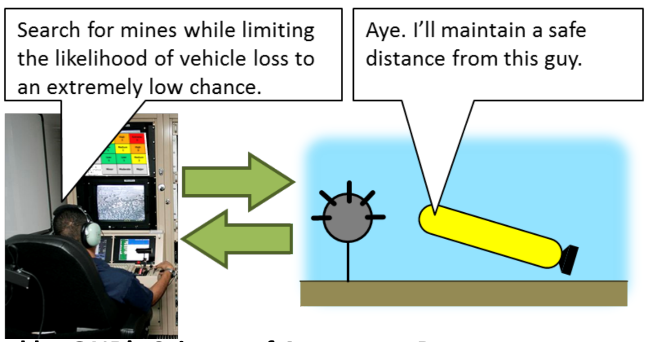

Work funded by ONR's Science of Autonomy Program

Technical Approach:

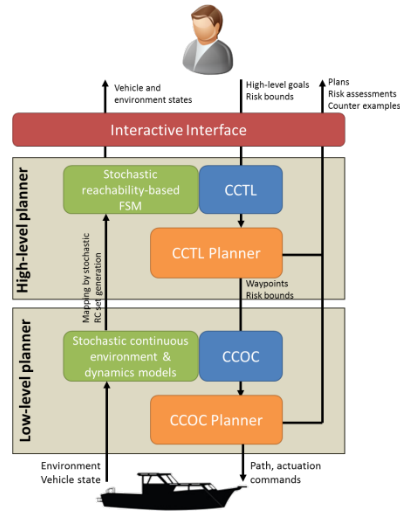

- Human specifies goals AND acceptable level of risk through an interactive interface
- Correctness guaranteed by chance-constrained temporal logic (CCTL) planner
- High-level state space representation is obtained through stochastic reachability/controllability (RC) set generation
- High-level plan encoded into a chance-constrained optimal control (CCOC) problem and executed by CCOC Planner

Why important to NASA:

- Exploring astrobiologically important targets (e.g., caves on Mars, subsurface ocean of icy moons) requires risk-aware on-board decision makings due to significant level of uncertainty and inability of orbital reconnaissance
- Future manned mission beyond LEO requires an Earth-independent system with a few astronauts cooperating with automated subsystems

Co-Investigators

- Prof. Missy Cummings (Duke)
- Prof. Behcet Acikmese (U. Texas at Austin)
- Prof. Ufuk Topcu (U. Texas at Austin)

RC-based spacecraft guidance

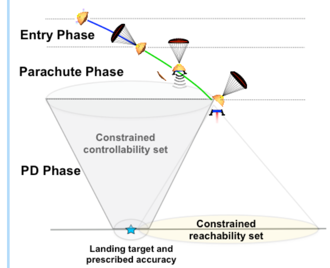

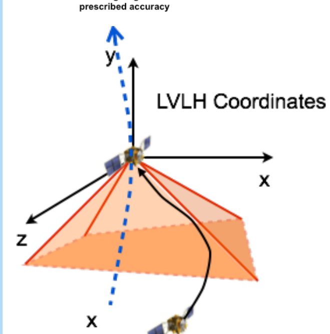

Human interface for monitoring uncertain system

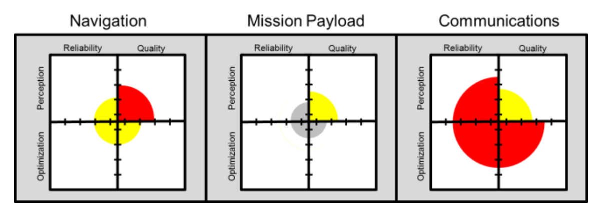

Safe Autonomous Rendezvous with CCOC

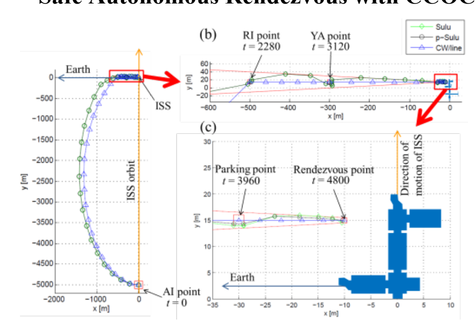

Risk-aware Personal Transportation System

Work performed at MIT with Prof. Brian Williams; Funded by Boeing

- **Goals**
  - Develop an autonomous, decentralized path-planning and scheduling algorithm that optimally balances safety and efficiency based on user's preferences
- **Approach**
  - Risk-constrained optimal path planning that allows users to specify a comfortable level of safety
    - Example 1: A risk-averse passenger flight maximizes safety by avoiding crowded/turbulent air space
    - Example 2: A time-critical military flight minimizes flight time by taking extra risk to go straight through crowded/turbulent air space
  - Market-based, decentralized air traffic control that guarantees collision avoidance and maximizes social welfare through bilateral, automated negotiations
    - A less time-critical flight makes way for a time-critical flight by receiving compensation
  - The proposed technology is built upon an existing risk-sensitive plan executive, *probabilistic-Sulu (p-Sulu)* [1], which has been developed by MIT/Boeing and demonstrated on MIT's SPHERES test bed [2]

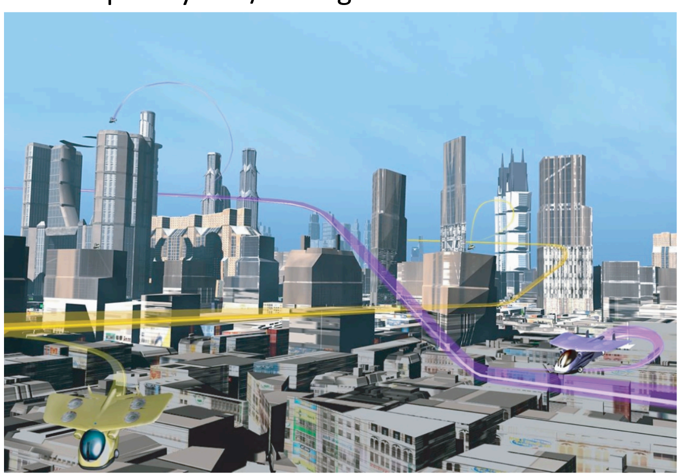

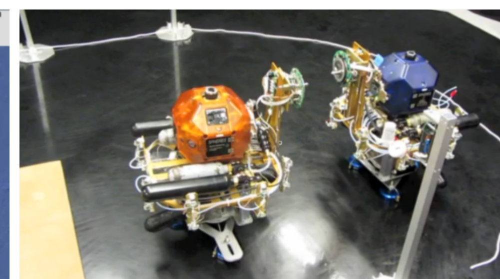

(Left) Risk-aware path planning for personal aerial vehicles by p-Sulu [1]
(Top) Demonstration of p-Sulu by MIT's SPHERES test bed [2]

[1] M. Ono, B. Williams, L. Blackmore. "Probabilistic Planning for Continuous Dynamic Systems under Bounded Risk." *Journal of Artificial Intelligence Research,* 2013
[2] C. Jewison, B. BcCarthy, D. Sternberg, C. Fang, D. Strawser. "Resource Aggregated Reconfigurable Control and Risk-Allocative Path Planning for On-orbit Assembly and Servicing of Satellites." Submitted to *AIAA-GNC*, 2014


Hidden Valley, Sol 712

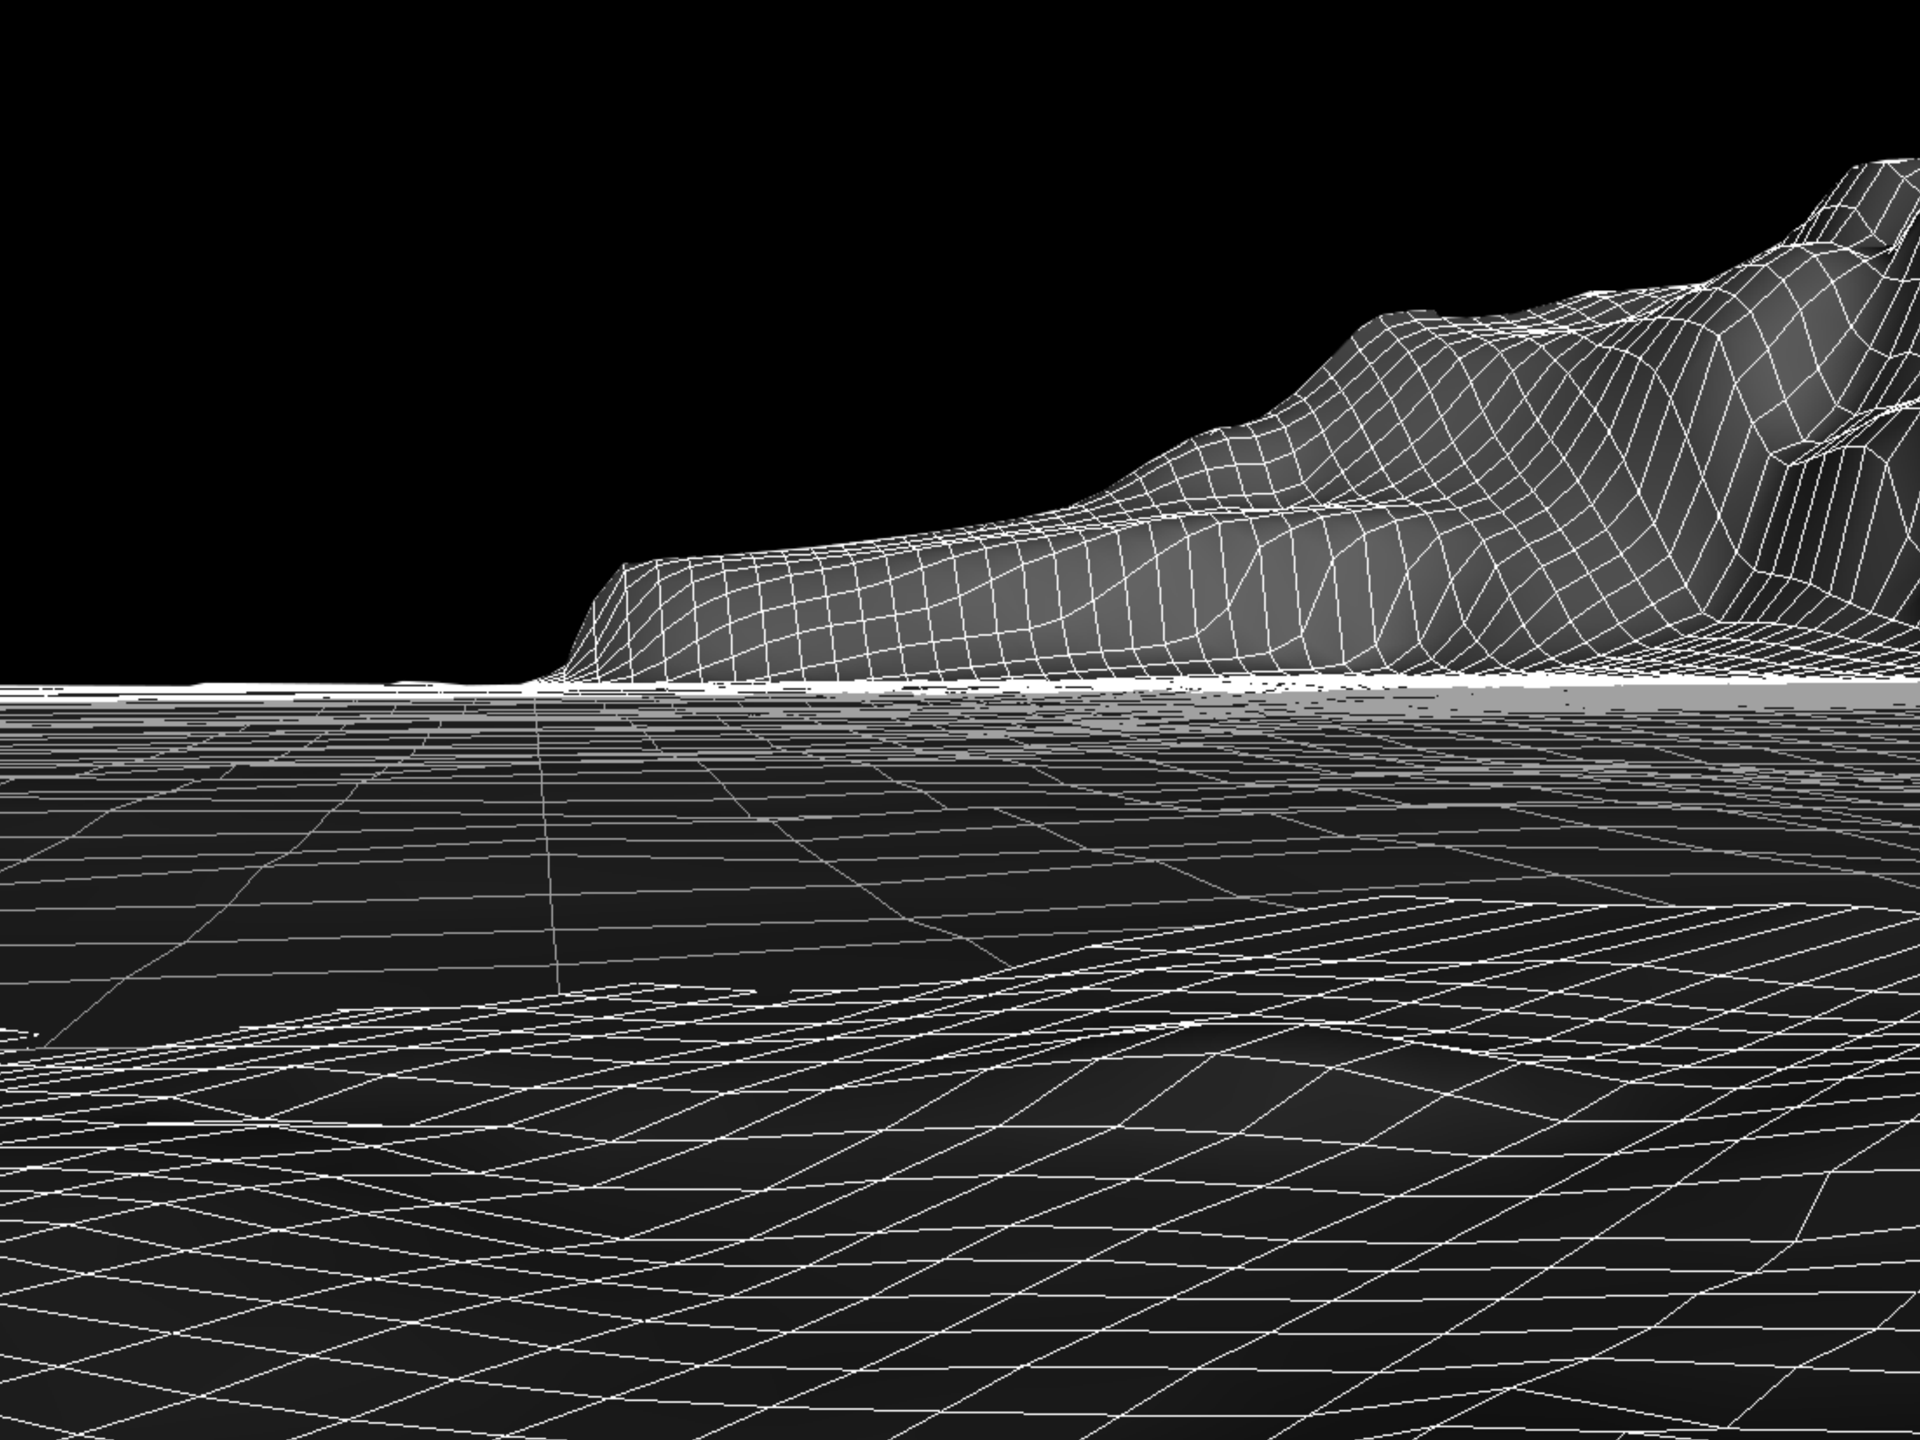

SPOC (Soil Property and Object Classification)

PI: Masahiro Ono (347E) ono@jpl.nasa.gov

- Visually classifies terrain types and features
- Based on deep convolutional neural network
- Trained by sparse labels generated by humas
- Classifies every pixel on images
- Demonstrated on a test rover in Mars Yard (off-board test)
- On-board variant being developed in RTD

SPOC-H: HiRISE deployment for M2020

Training Data Output

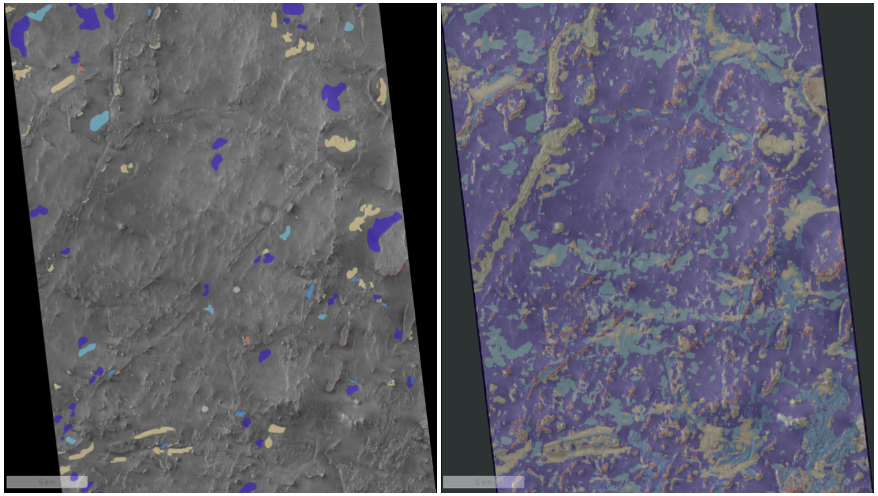

Smooth regolith Smooth outcrop Fractured outcrop Sparse linear ripples Rough outcr
Crater Rock field Dense linear ripples Polygonal ripples Deep sand Scarps

Work by Brandon Rothrock

SPOC-G: NAVCAM deployment for MSL

Dingo Gap

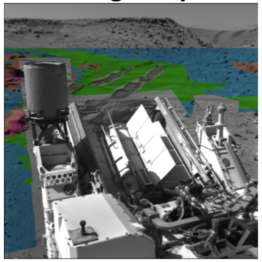

Hidden Valley

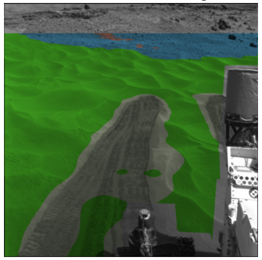

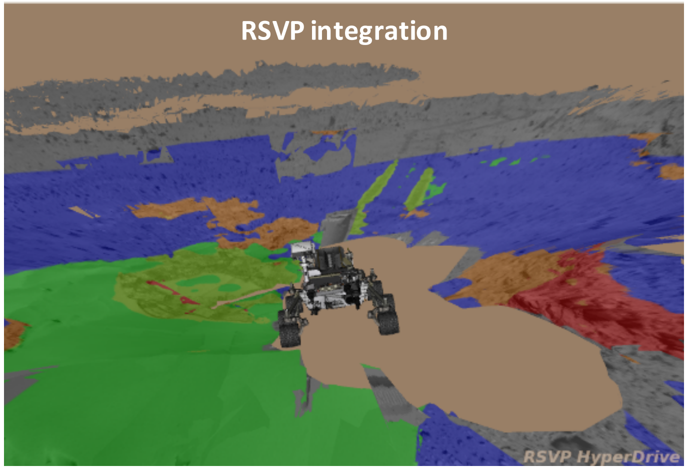

Sand Smooth Outcrop
Rock field Outcrop w/ rocks

Work by Ryan Kennedy

Mars Yard Demo

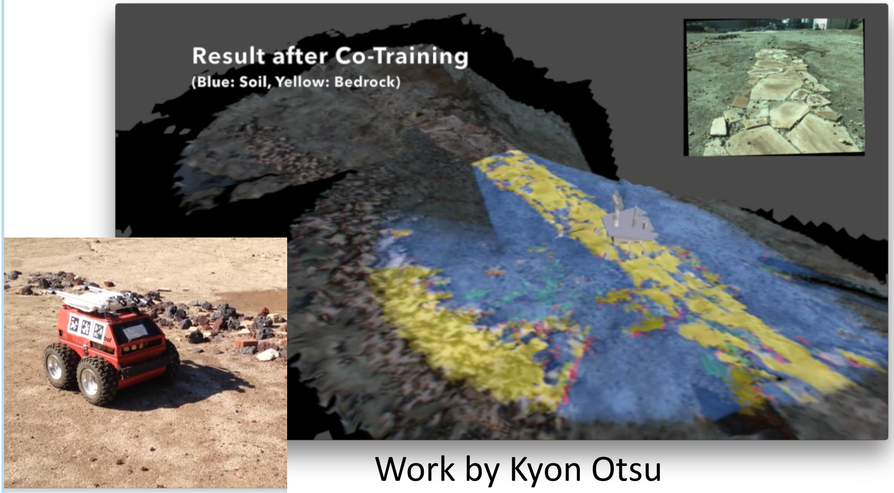

Work by Kyon Otsu

Hidden Terrain Inference with Gaussian Process

Work by Clark Zhang (PhD student @ UPenn)

- On-line learning of kernel to capture the "pattern" of terrain
- Gaussian process regression with the learned kernel to estimate the occluded terrain
- Use quad-tree to partition the state space with varying resolution

Mars terrain

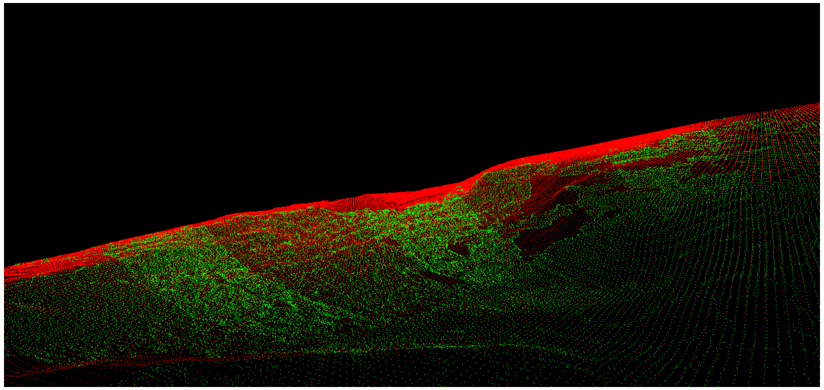

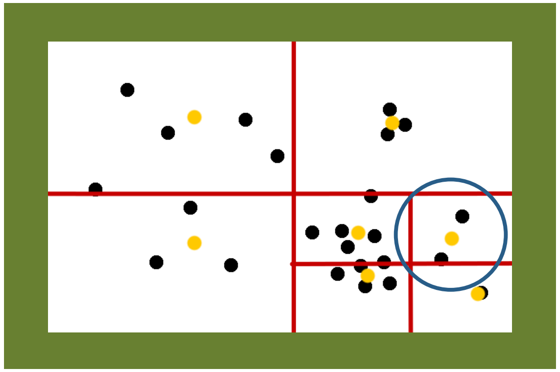

Volcanic fissure

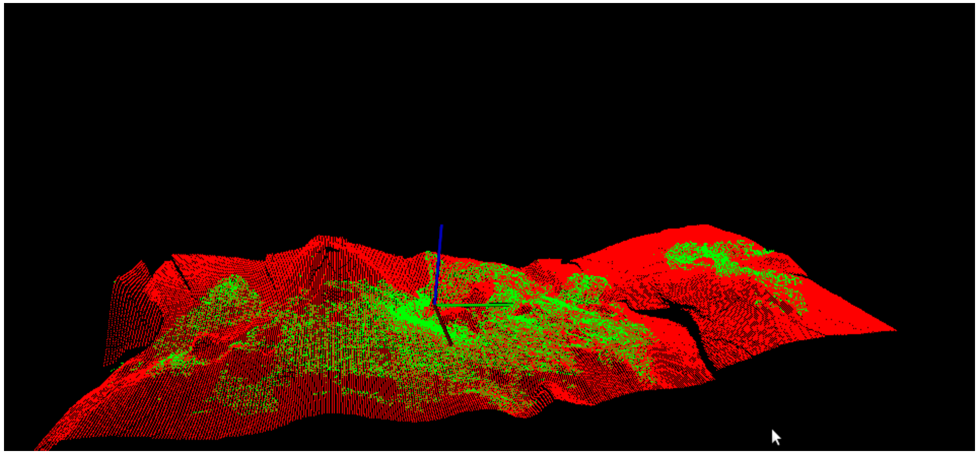

Advanced Path and Motion Planning

Masahiro Ono (347E) ono@jpl.nasa.gov

- Safe path/motion planning algorithms for environments with significant risk
- Motion planning with wheel placement for navigating through highly cluttered environment
- Sequential Dijkstra Algorithm for co-optimizing route and goals
- Planning with spacio-temporally extended graph for cooperative multi-vehicle path planning

Co-optimization of goal and route for M2020 landing site traversability analysis

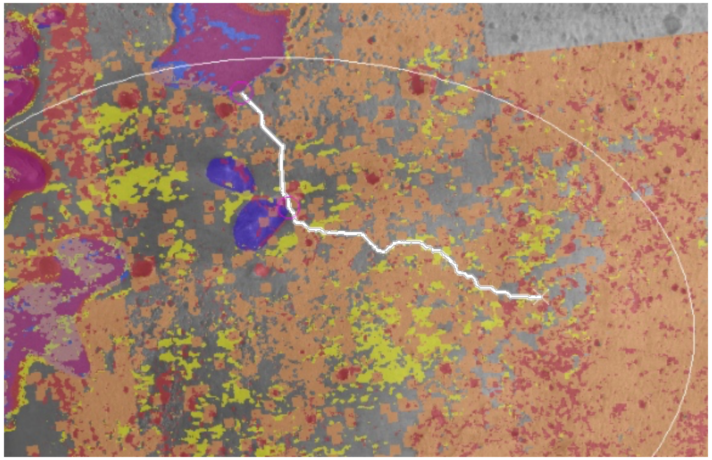

Multi-vehicle coordination

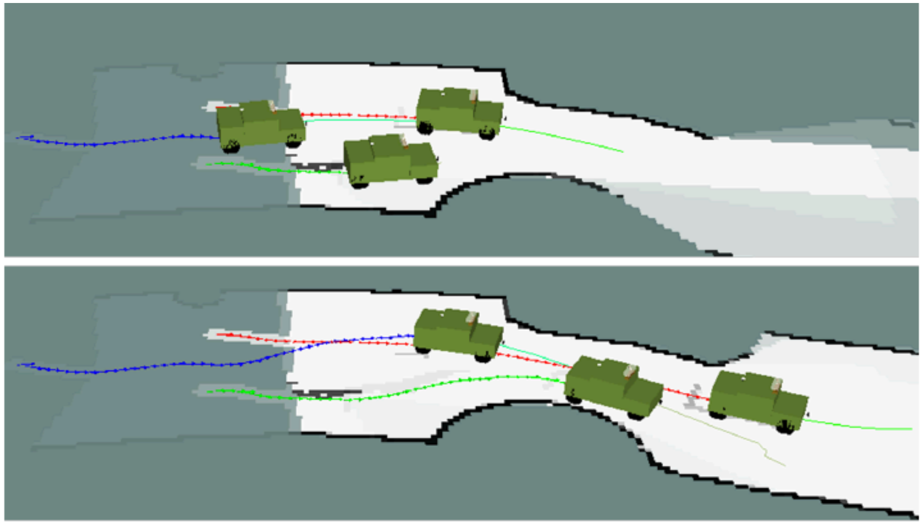

Work with Greg Droge, Amir Rahmani, et al

Motion planning for cluttered environment

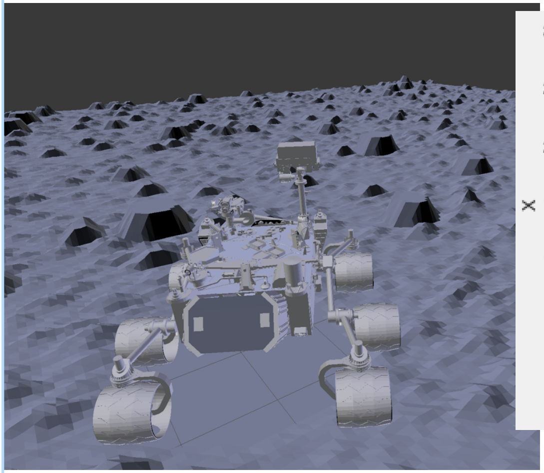

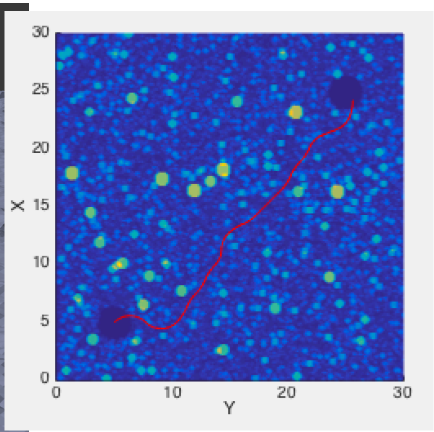

Work by Kyohei Otsu

MTTTT (Mars 2020 Traversability Tools)

M. Ono, B. Rothrock, E. Almeida, H. Gengle, F. Calef, T. Solomon, R. Otero, A. Huertas, A. Ansar, M. Heverly

- Supports M2020 landing site selection
- Evaluates terrain type on HiRISE by SPOC
- Evaluates rock abundance through automated rock detection
- Optimizes the combination of ROIs (regions of interest) to visit and route by Sequential Dijkstra Algorithm
- Performs Monte Carlo simulation to obtain the probability distribution of driving distance and time to satisfy the mission requirements

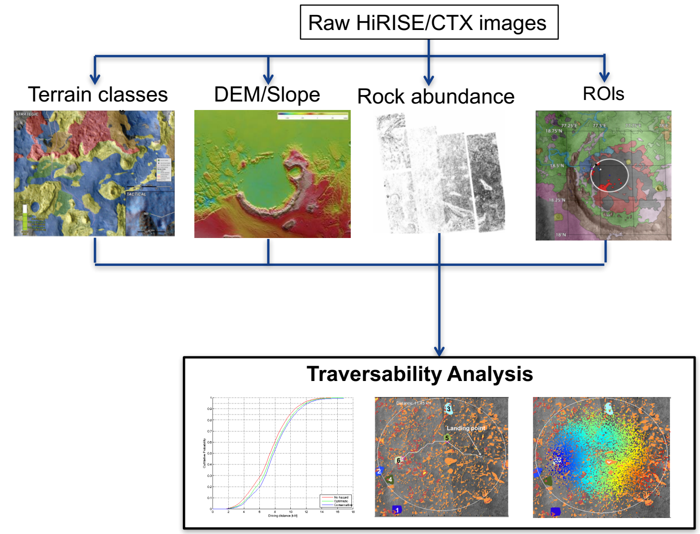

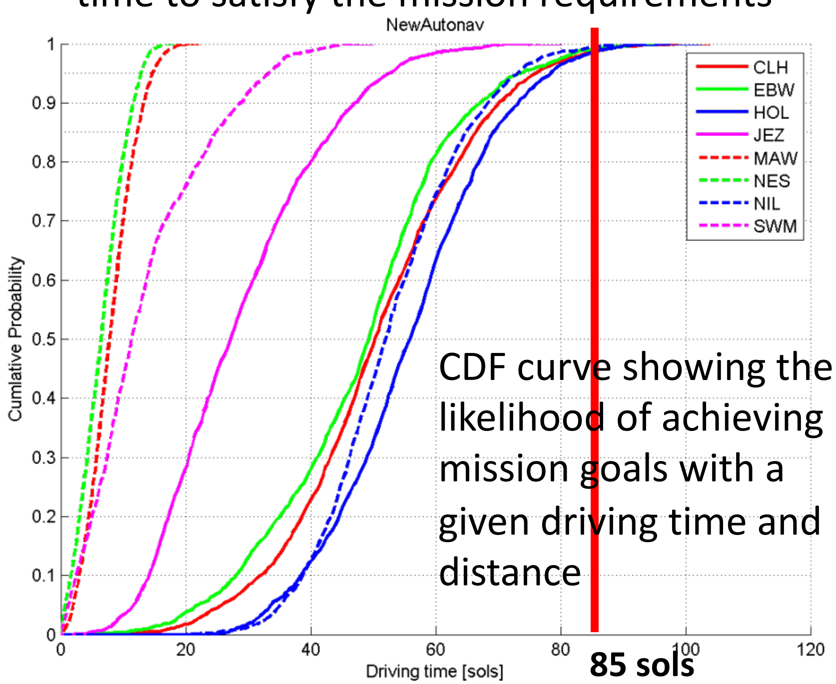

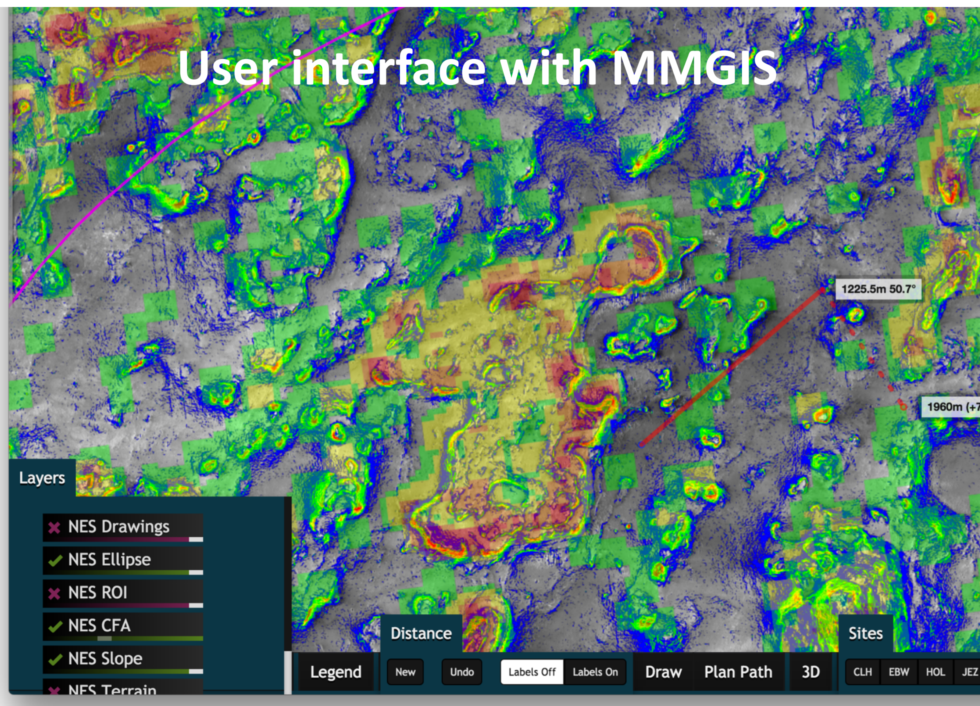

Next-Gen Autonomy for Mars Rovers

PI: Masahiro Ono (347E) ono@jpl.nasa.gov

Project objective:

Enable Autonav (autonomous navigation) to *safely* traverse over more difficult terrain (>10% CFA rock abundance, mixed terrain types) at a faster rate (>80 m/hr, ~5x faster than MSL)

Technical Approach:

- Situation awareness: risk evaluation based on visual recognition of terrain type and topology (e.g., avoid >5 deg slope on sand dunes)
- Advanced motion planning for cluttered environment: remove conservatism by straddling over small rocks
- Extensive V&V on test rovers in JPL's Mars Yard

***Work funded by JPL's R&TD program**

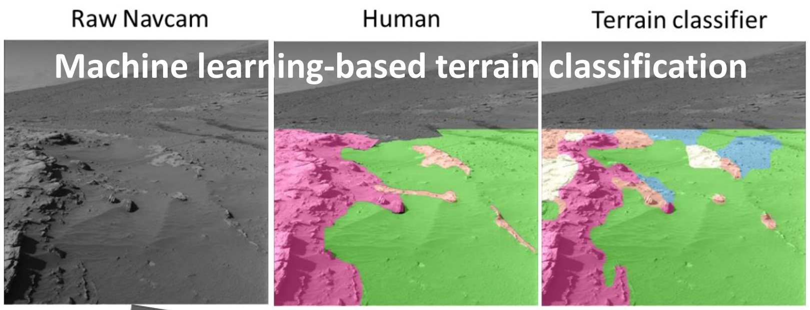

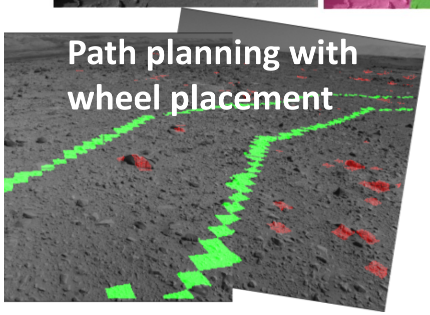

Why important to NASA:

- Scientifically interesting sites tend to have high rock abundance and complex topology (e.g., Jezero Crater, NE Syrtis, Melas Chasma)
- M2020 and Sample Retrieval and Launch rovers would benefit from Autonav because they are required to traverse 60-80 m/hr *on average*
- Departure from the current costly, Earth-dependent operation model (e.g., ~$30M per rover per year for MSL) is necessary to expand the scale of surface operation
    - Future human mission would be likely to operate multiple ISRU rovers before the arrival of astronauts

FY15 Mars Yard Demo (Video)
(Terrain classification w/ co-training)

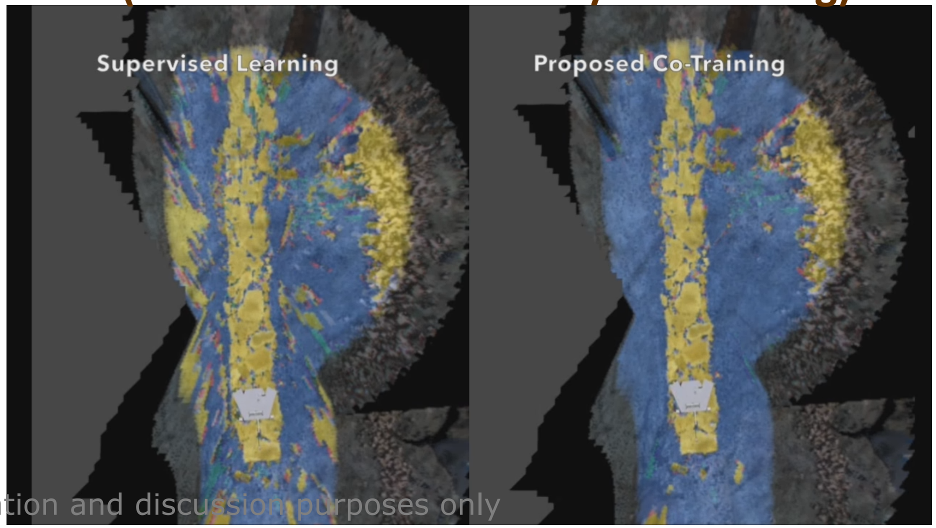

Journey to the Center of Icy Moons

Masahiro Ono(PI), Mitch Ingham, Tara Estlin, Aaron Parness, Karl Mitchell, Kevin Hand
Jet Propulsion Laboratory, California Institute of Technology

Concept: *Icy-moon Cryovolcano Explorer (ICE)*

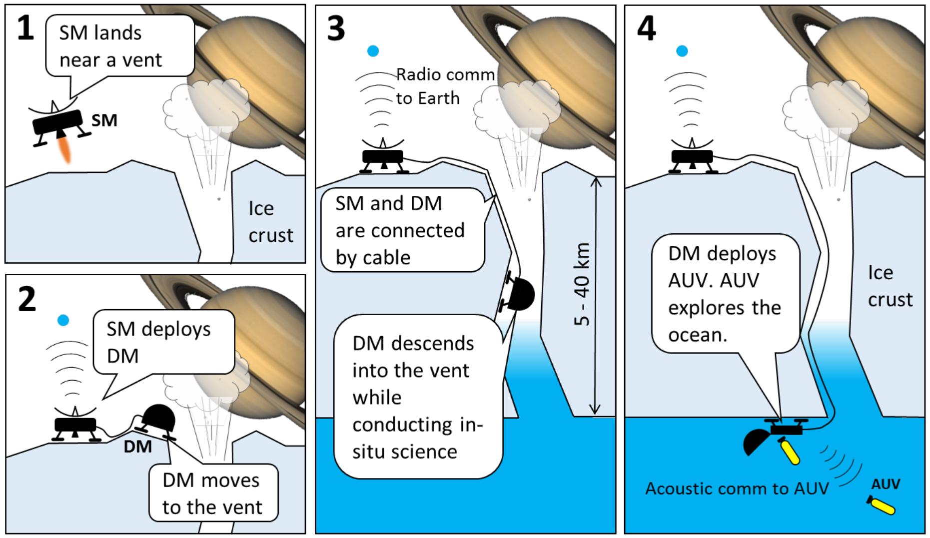

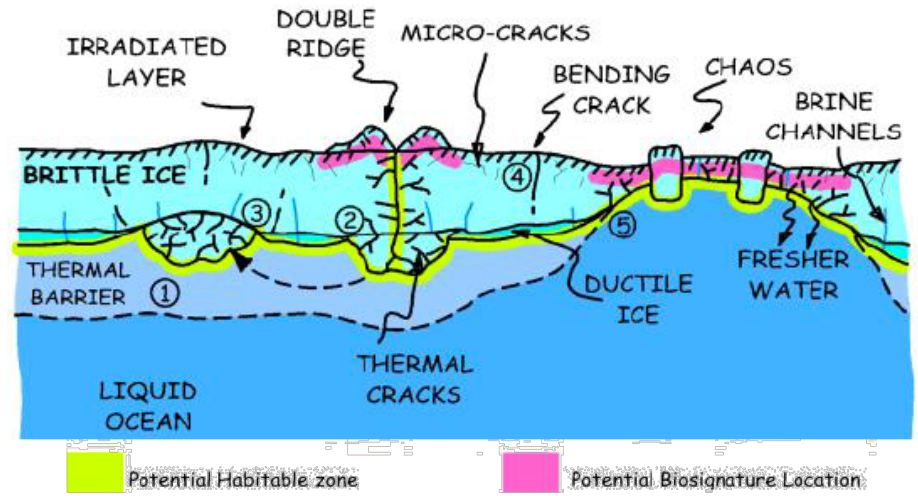

Vent is not only a gateway to the subsurface ocean, but also a potential habitable zone. (Figueredo et al. 2003)

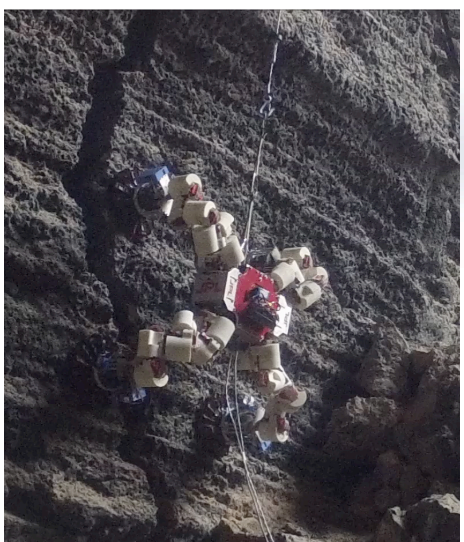

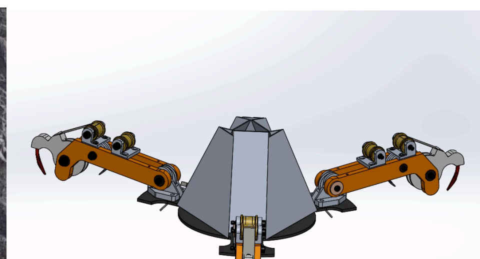

↑Proposed ice gripper with claws (Carpenter)
←LEMUR climbing robot (Parness et al)

Risk-aware autonomy demonstrated on a rover

Risk-aware Machine Learning for Resilient Space Exploration

Co-PIs: Yisong Yue and Hiro Ono

- Big goal: paradigm shift in robotic space exploration - from robust to adaptive
- Challenges:
  - Unpredictability in robot's behavior: how to guarantee safety?
  - Machine learning needs a lot of data, but robot has to adapt to new environment quickly (e.g., Venus lander would last only for a few hours)

Safe exploration

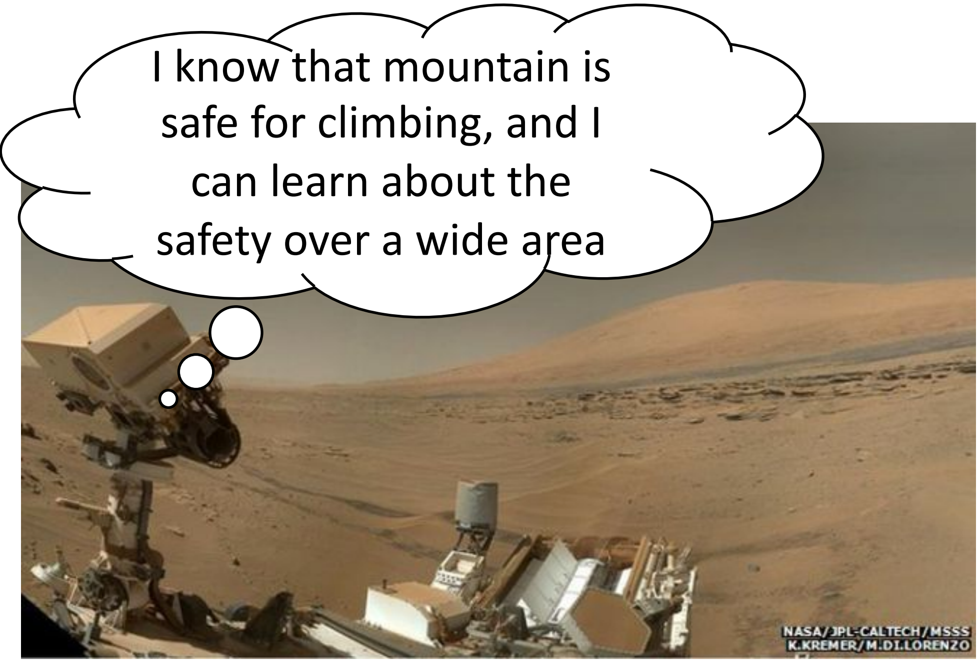

Hybrid Model-Based and Model-Free Planning

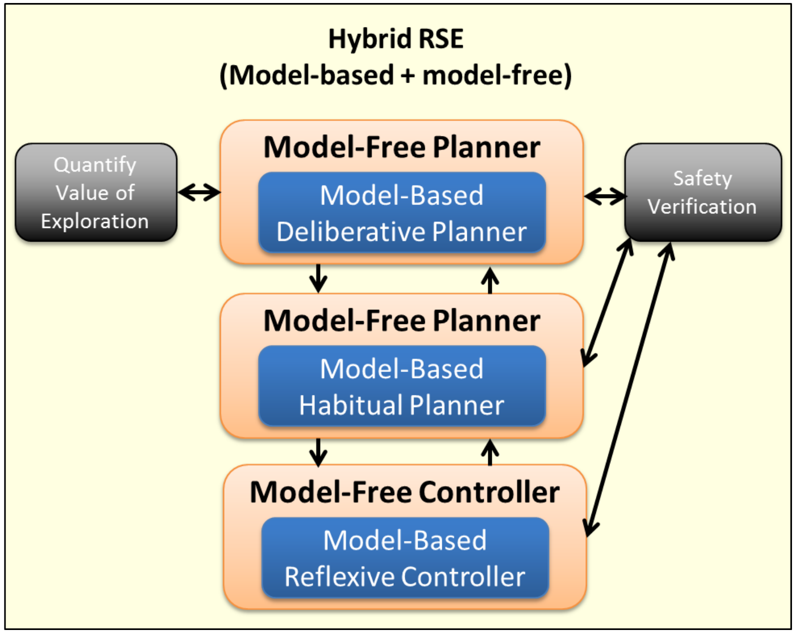

Backups

Comet Hitchhiker

FY15 NIAC Phase I Study

Masahiro Ono
Jet Propulsion Lab, Caltech

Paper in the proceedings:
"The Hitchhiker's Guide to the Outer Solar System"

- Hitches rides on small bodies by using a harpoon and an extendable momentum exchange tether
- Would enable rendezvous with small bodies in the outer Solar System (e.g., KBOs, Centaurs)
- 1.5 km/s hitchhike feasible with Zylon tether; 10 km/s enabled by CNT tether

SMART

Space Mission Architecture and Risk Analysis Tool

Masahiro Ono

Objective

Support high-level trade study for a *multi-mission campaign,* such as Mars Sample Return

- Perform trade-offs between success probability, utility (e.g., science return), and cost

Deliver capabilities to

- Capture high-level problem structures while providing an interface to detailed analysis
- Marry intuitive/qualitative arguments with mathematical rigor
- Discovers hidden/unintuitive causal links between seemingly isolated factors

System Architecture

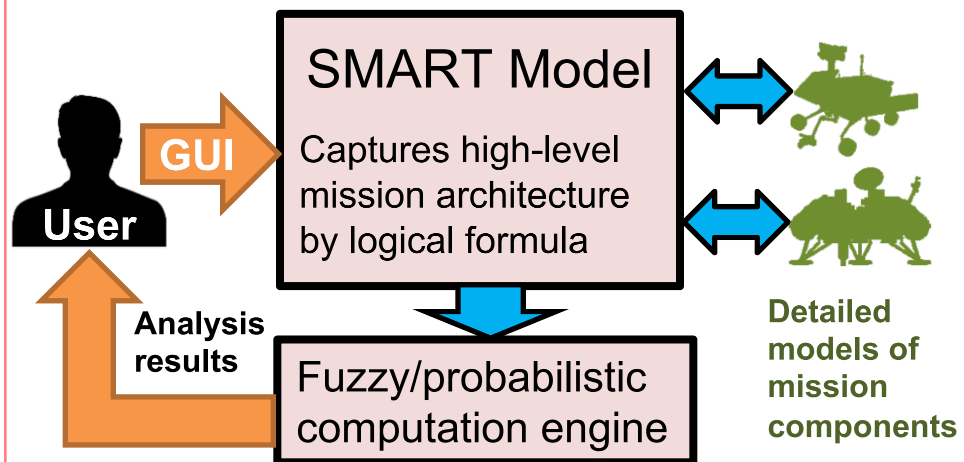

Approach

- Represents dependencies between design params and mission outcomes by logical formula
- Supports fuzzy and probabilistic computation
- Provides Simulink-based GUI to build models
- Performs symbolic computation of Boolean algebra
- Any fuzzy operators (e.g., Zadeh operators) as well as probability operation can be used

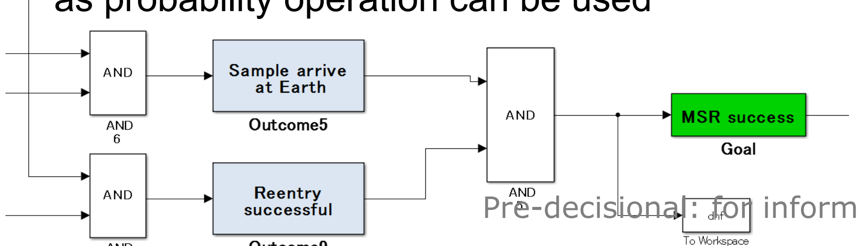

Analysis example on Mars Sample Return

SMART Model of MSR

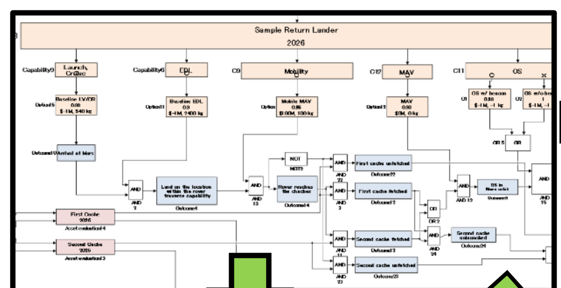

Sensitivity analysis

```
Sensitivity Analysis Result
pLV/CR2 : 0.3865
pLV/CR1 : 0.35575
pLV/CR3 : 0.35575
pMOB1 : 0.35454
```

1% increase in the reliability of rover mobility results in 0.35% increase of MSR success probability

Monte Carlo analysis

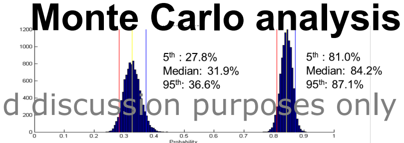

Risk assessment

CEMAT: Risk-based EDL & mobility analysis

- Combined EDL-Mobility Analysis Trade Study Tool
- Optimizes landing ellipse, surface path, and EDL control policy within a given bound on the probability of landing failure

Input

- Terrain (hazard/rock/slope maps)
- Science targets
- Lander/rover design params (control authority, rover size, etc)
- Bound on the probability of landing failure (e.g., 1%)

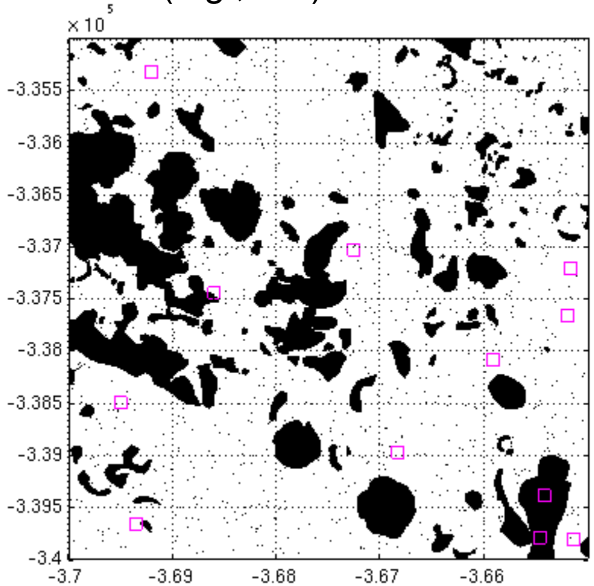

CEMAT

Output

- Expected driving distance to visit a specified number of science target
- Resulting probability of landing failure
- Landing ellipse
- Optimal surface path
- EDL control policy