




Distributed Spacecraft with Heuristic Intelligence to Enable Logistical Decisions for Global Navigation Satellite System Reflectometry (GNSS-R) of Wildfires

Presenter: Vinay Ravindra^{[1][2]}
vinay.ravindra@nasa.gov

Team: Sreeja Roy Singh^{[1][2]} (PI), Mahta Moghaddam^[3], Kurtis Nelson^[4], Richard Levinson^[1], Amer Melebari^[3], Archana Kannan^[3]

1. NASA Ames Research Center, Moffett Field, CA
2. Bay Area Environmental Research Institute, Moffett Field, CA
3. University of Southern California, Los Angeles, CA
4. USGS EROS Center, Sioux Falls, SD



Adaptive, Intelligent and Responsive New Observing Strategy (AIR NOS) for Wildfires

Adaptive

Science drivers dictating observations, based on previous observations

- Geophysical retrievals
- Data assimilation into NWP
- Value framework, forecast reporter

Intelligent

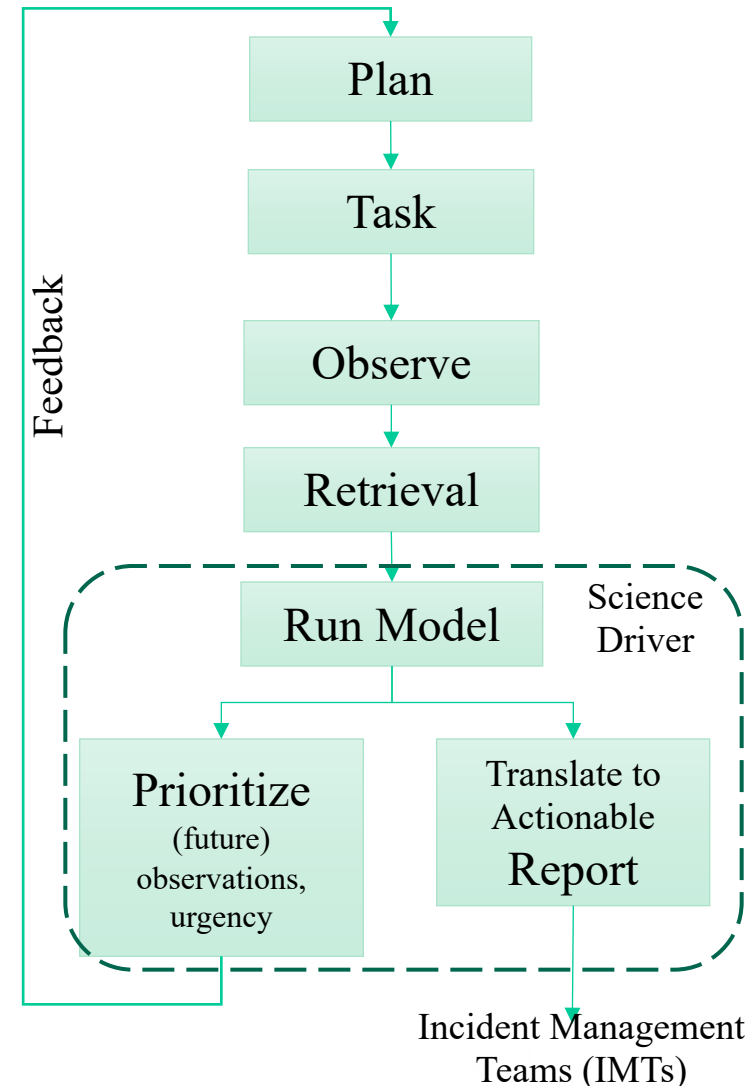
Efficient resource utilization using AI

- Planning, orbit/coverage calculations, satellite sub-system modeling, downlink data prioritization

Responsive

Time of event to delivery of report

- Use of only NRT data, cognizance of processing times, sync of operations



Global Navigation Satellite Systems Reflectometry (GNSS-R)

- L-band (microwave) observations, can penetrate clouds, smoke
- Observations of specular reflections => complimentary, **unique data** compared to backscatter (SAR), radiances (radiometer)
- Passive bistatic radar => small form factor of satellites => numerous distributed satellites => **high temporal sampling**
 - NASA Cyclone Global Navigation Satellite System (CYGNSS) mission
- Has demonstrated scientific retrieval of surface, including wind speed over ocean, and soil moisture and flooding over land.




Fig: GNSS-R imaging geometry

Rose, Randy, Scott Gleason, and Chris Ruf. "The NASA CYGNSS mission: a pathfinder for GNSS scatterometry remote sensing applications." *Remote Sensing of the Ocean, Sea Ice, Coastal Waters, and Large Water Regions 2014*. Vol. 9240. SPIE, 2014.




Fig: Delay Doppler Map (DDM)

Adaptation of scientific and technological benefits offered by GNSS-R to the wildfire application.

Pre-fire mission

Prioritize observations over areas where there is greater chances of fire.

- Use of USGS Fire Danger maps as science driver
- Retrieve from soil moisture derived from (uncompressed) GNSS-R
- Feedback to refine fire-danger predictions

Active fire mission

- Prioritize fast delivery of observations at active fire area
- Retrieve burnt area locations from GNSS-R
- Assimilate into WRF-SFIRE

Default Mission

- Model of 'nominal' mission operations

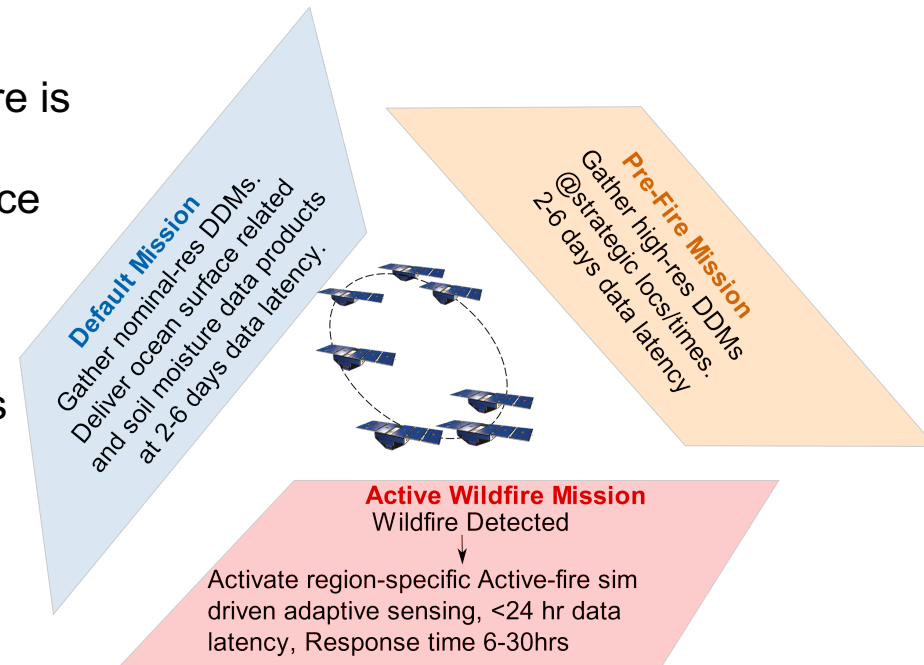


Fig: Simultaneous execution of multiple missions



Preliminary work

6

Independent development of components of the proposed adaptive sensing framework

- Enhancing USGS Fire Danger product with GNSS-R (CYGNSS) derived soil moisture
 - Goal: Improved Fire Danger by adjusting Fuel Moisture with observed data
- Exploration of physics-based and ML-based detection of 'burnt-area' with GNSS-R (CYGNSS) Delay Doppler Maps
 - Goal: High cadence 'burnt-area' detections during active fire
- Modeling & Intelligent tasking algorithm development & simulated tests with CYGNSS as example
 - Goal: Efficient use of satellite resources for maximizing science value gathering with low latency



Enhanced Fire Danger product

7

- USGS Fire Danger Forecasting produces 7-day forecast products for fire potential index, large fire probability, and fire spread probability.
- Utilizes a combination of satellite-derived vegetation indices, various biogeophysical variables, and weather information.
- Modify the nominal workflow of calculating FPI to consider observed **soil moisture anomalies**
- 10hr FM term is adjusted by the anomaly term

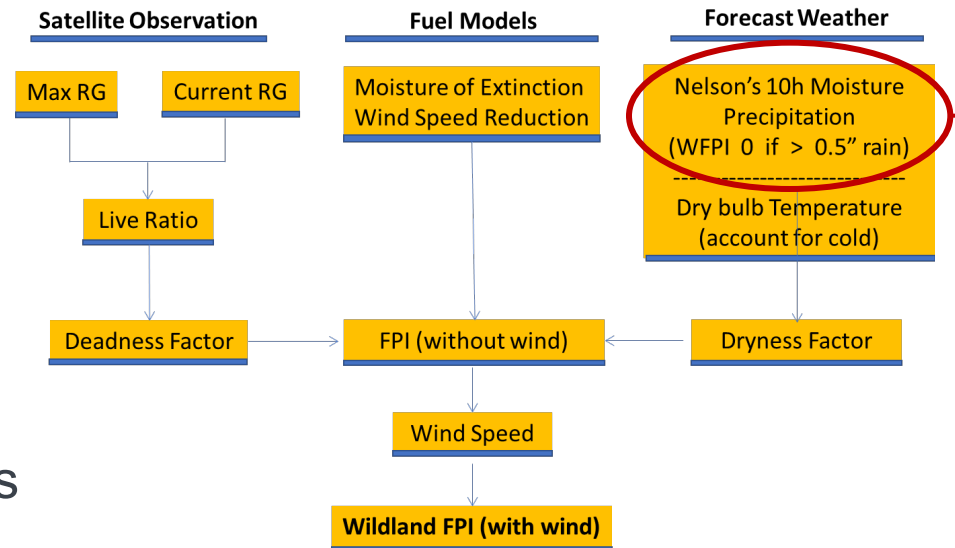


Fig: Fire Potential Index Data Flow

$$\%SM_{rel} = \frac{SM - SM_{avg}}{SM_{avg}}$$

CYGNSS 9km Soil Moisture products used for adjustment

Validation

- Three test cases from MTBS record:
 - Vivian Fire – 11825 acres, ignition 8/20/2019
 - Game Ranch Fire – 3022 acres, ignition 7/13/2020
 - Gate 5 Fire – 11457 acres, ignition 7/14/2020
- Adjusted WFPI is **higher at locations which later caught fire**
 - Mean difference is +13

Operational WFPI



Adjusted WFPI



Fig: Grayscale image of operational and adjusted WFPI (unit-less number that ranges from 0 to 150)




Fig: Histogram of the adjustment at the test events

Next Steps:

Incorporation of quality checks and testing of with more past fires.

Physics based detection

- Improved geometric optics with topography (IGOT) model compute GNSS-R DDM

- Retrieve vegetation parameters, by matching forward-model (IGOT) output with actual observation.

- Change of vegetation characteristics **between DDMs before the fire and after** is used to for detection.

- SMAP is used for surface soil moisture.




Fig: Physics based retrieval

Optimize to find vegetation params of the IGOT model which produce a matching DDM


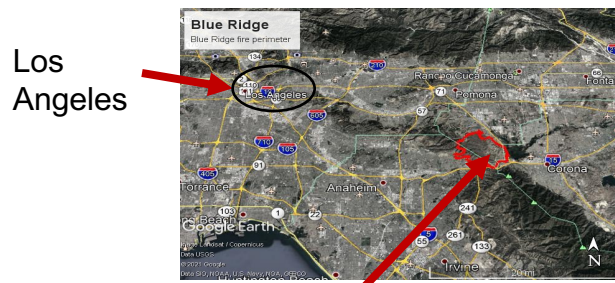


Fig: Observed CYGNSS DDM

Fig: Simulated IGOT DDM




Test case: *Blue Ridge fire*

- Duration: Oct 26, 2020 – Nov 7, 2020 (11 days)
- Burned area: 13,964 Acres




Fire perimeter

Cannot classify based on
vegetation parameter retrievals
or specular point reflectivity



Good matching
between CYGNSS
and forward model
specular point
reflectivity [RMSE:
0.5 dB]

Next Step:
Retrieval of *both* Soil
Moisture & Veg
params with *multiple*
DDMs.



GNSS-R Burnt Area detection

11

ML based detection

- Random Forest Model

Inputs

- CYGNSS DDM: Specular point incidence angle, SNR, Reflectivity of specular point

Ancillary data


- SRTM: Elevation, Slope, Aspect
- SMAP: Soil Moisture, Vegetation Water Content, Surface Temperature, Precipitation

Truth data

Overlapping dataset created with:

- Landsat Collection 2 Level-3 Burned Area Science Product,
- MODIS Burned Area,
- ESA CCI Burned Area.

Test case: Blue Ridge Fire



Train Confusion Matrix

Test Confusion Matrix

	Train (70%)	Test (30%)
Accuracy	100 %	97.5 %
F1 score	1.0	0.8717

- The burned area data is highly imbalanced.
- Good accuracy and F1 score improvement with important features (Reflectivity, SNR, incidence angle, Soil moisture, VWC, surface temperature)

Next Work: Explore more locations with less complex topography.

While LEO small satellites enable distributed, high-temporal sampling, they are highly constrained

- Onboard data, downlink rate, power, re-orientation speed, etc.

Numerous options

- Observation modes, Selection of target, Downlink priority

Multiple objectives

- How to balance priorities of Pre-fire, Active-fire, Default missions?
- Different missions shall have different
 - Pre-fire: Data latencies can be relaxed
 - Active-fire: Data latency is critical

Some of the agility features

Mode	Integration time (s)	Bit depth	# specular locations	Data Volume increment factor
Nominal	0.5 (2 Hz)	9	4 (max SNR)	1
Modified	0.25 (4 Hz)	9	4 (max SNR)	2
Full DDM	0.5 (2 Hz)	32	4 (max SNR)	~48.7
RawIF	N/A	32	Unrestricted (L1 band)	6724.30

Table: Operational modes of observation available on CYGNSS




Fig: Prioritized downlink from onboard data buffer

Each choice comes with an associated reward/ cost which can be modeled with satellite orbit, subsystem (data, power) dynamics. Coupled with intelligent tasking, we can operate efficiently.



Monte Carlo Tree Search (MCTS), Coulom 2006

- First algorithm to beat a human champion in the game 'Go'.
- Used by Tesla cars to avoid collisions
- Good for extremely large search spaces (decision spaces)
- Integrates Planning with Reinforcement Learning (RL)
- Explores search spaces using Monte Carlo simulations ("rollouts")
- Collects outcomes from each rollout to update reinforcement learning statistics
- Each rollout is a training case for reinforcement learning
 - Keeps statistics on *expected reward* for different choices
- Search involves balancing choices between *Exploration* vs. *Exploitation*
 - Exploration: Prefer new choices *which have not yet been simulated*
 - Exploitation: Prefer choices which worked best in prior simulations (highest expected reward based on RL)
- "Anytime Algorithm"
 - Planning can be interrupted at "any time" and a valid solution will be returned
 - More planning time is guaranteed to monotonically improve solution
- Specialized D-SHIELD adaptations to optimize satellite observation planning

Rémi Coulom (2006). "Efficient Selectivity and Backup Operators in Monte-Carlo Tree Search". *Computers and Games, 5th International Conference, CG 2006, Turin, Italy, May 29–31, 2006.*

Feature important for responsive missions.



Simulation Scenario:

Pre-Fire Mission (1st Aug 2020, 1 day)

- 1 CYGNSS satellite
- 3 ground-stations @ HI, Chile & W. Aus.
- **RawIF mode** (high data volume) observation over locations of high fire danger
- FIFO data buffer

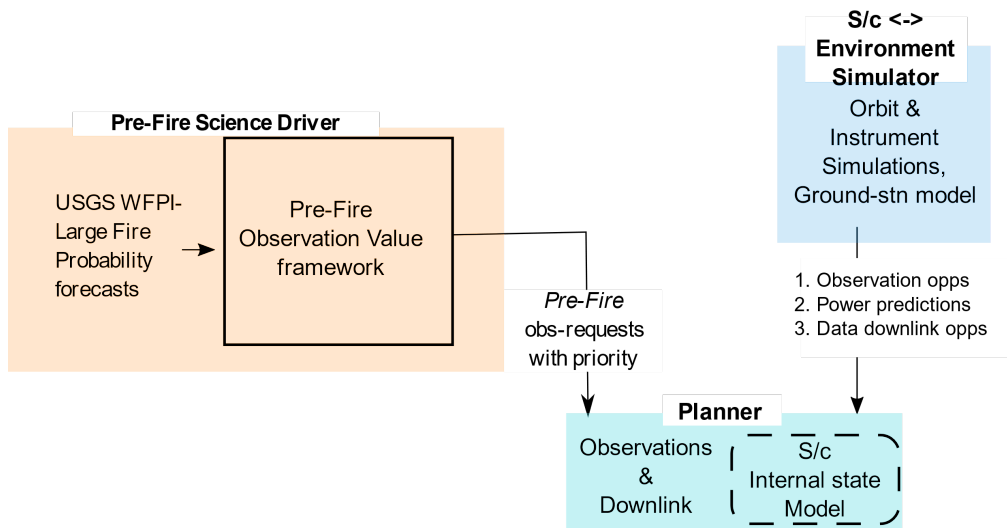


Fig: Pre-fire mission execution flow

Objective formulation

Maximize the cumulative sum of WLFPs of the **imaged** & **downlinked** locations.

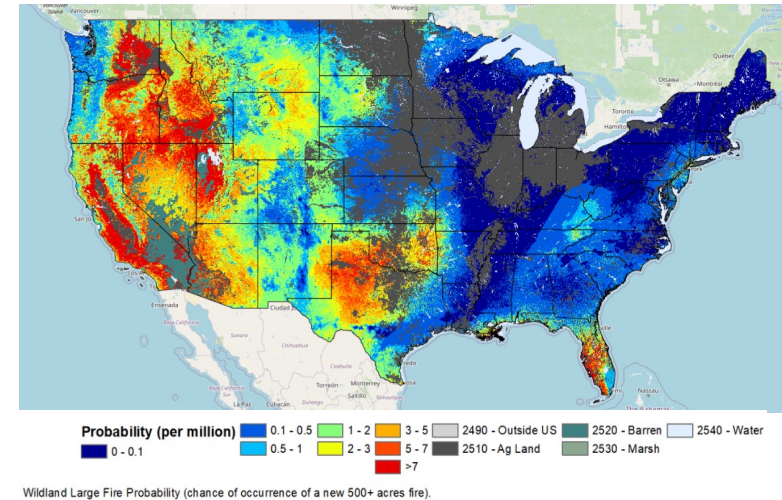


Fig: USGS WFPI-based Large Fire Probability (WLFP) as science driver



Orbits Access

- GNSS-R specular coverage calculation with multiple GNSS satellites


Data

- Max capacity ~ **5.7 Gb**, RawIF mode data rate = **96.22 Mb/s**, Downlink rate = **4 Mb/s**

Constrained Onboard data storage, downlink rate, and available power

Power

- Max charge: 86.4 Watt-hours, Avg Solar power in: **70 W**, Bus + sensor power consumption: **38.3 W**, Downlink power consumption: **22.6 W**, Minimum charge: 55% of max charge



Intelligent Tasking & Modeling

16

D-SHIELD Fire search space is immense

For each satellite in a 24-hour period:

~ **10,000 decision variables**

Each decision variable represents 1 second when the planner must make a **binary choice**:

- Make an observation or not
- Downlink data or not

The # of states to explore is:

$2^{E10000} =$
Infinity

(Google Exponent
Calculator)

Snapshots of the **Best Plan** after 200k rollouts (15hr process time)

Idling

```
time: 1, ***, bat. 89.99 % (eclipse)
time: 2, ***, bat. 89.98 % (eclipse)
time: 3, ***, bat. 89.96 % (eclipse)
time: 4, ***, bat. 89.95 % (eclipse)
```

Downlink

```
time: 22209, ***, bat. 83.3 % (eclipse)
time: 22210, ***, bat. 83.29 % (eclipse)
time: 22211, DNL, bat. 83.27 %, storage: 4586.568, gpCount: 0, score: 1398.514, targets: CHI (eclipse)
time: 22212, DNL, bat. 83.26 %, storage: 4582.568, gpCount: 0, score: 1398.903, targets: CHI (eclipse)
time: 22213, DNL, bat. 83.24 %, storage: 4578.568, gpCount: 0, score: 1399.292, targets: CHI (eclipse)
time: 22214, DNL, bat. 83.22 %, storage: 4574.568, gpCount: 0, score: 1399.686, targets: CHI (eclipse)
time: 22215, DNL, bat. 83.2 %, storage: 4570.568, gpCount: 0, score: 1400.08, targets: CHI (eclipse)
```

Imaging

```
time: 3074, RAW, bat. 99.99 %, storage: 2694.076, gpCount: 0, score: 289.817, targets: [24688, 25019, 27928, 31294, 33800, 34085, 35231, 42057, 42287]
time: 3075, IDL, bat. 99.99 %
time: 3076, RAW, bat. 99.99 %, storage: 2790.293, gpCount: 0, score: 305.391, targets: [24689, 25020, 27929, 31001, 31295, 31593, 31870, 33525, 34080, 35232, 35502, 35507, 35770]
time: 3077, RAW, bat. 99.99 %, storage: 2886.51, gpCount: 0, score: 318.177, targets: [25021, 27606, 31296, 34086, 34354, 35503, 42059, 42289]
time: 3078, IDL, bat. 99.99 %
```


Animation of **Best Plan** after 200k rollouts (15hr process time)

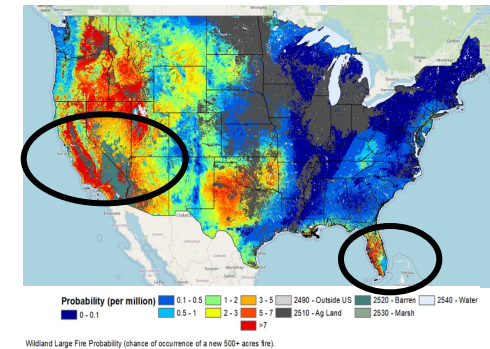
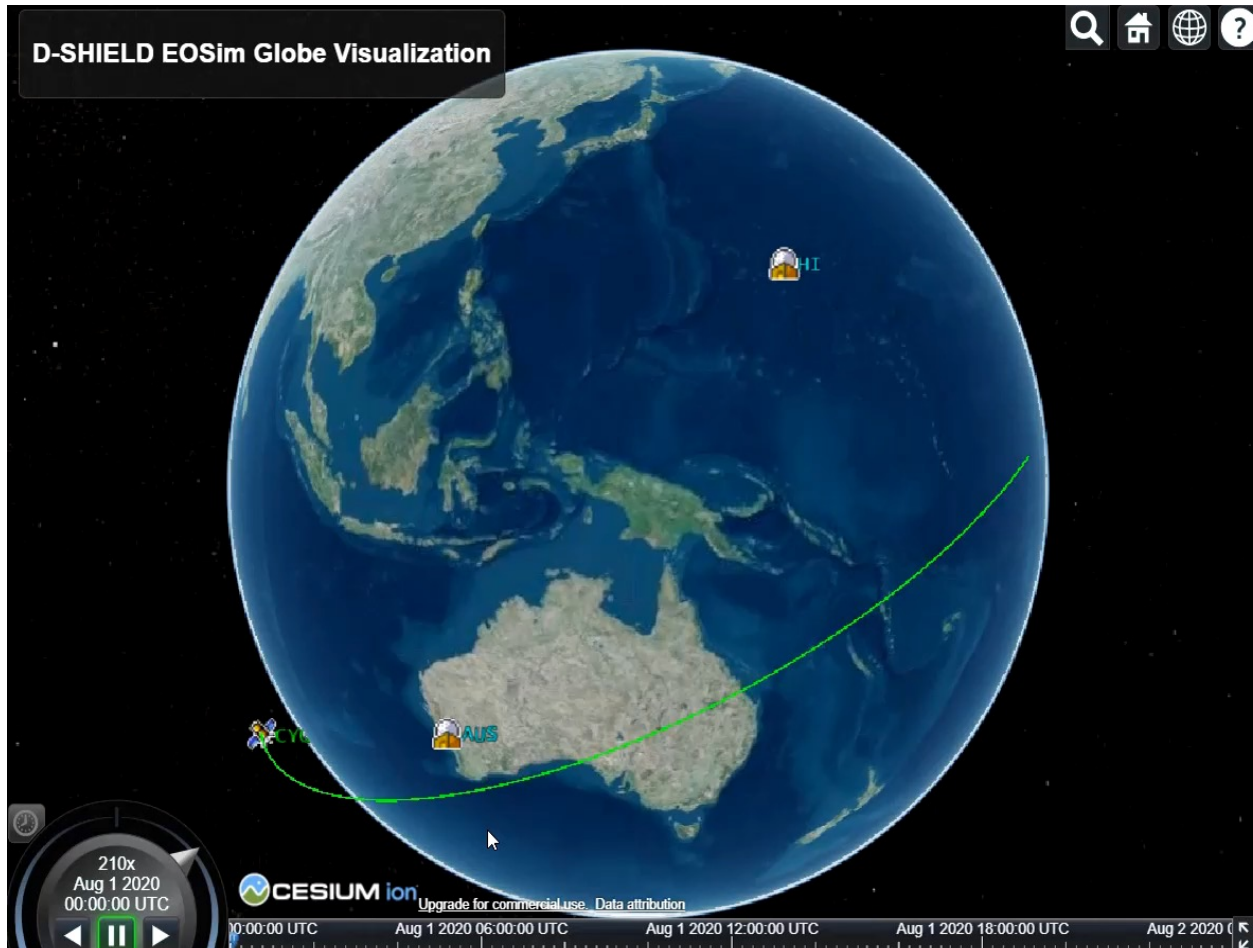



Fig: USGS WFPI-based Large Fire Probability (WLFP) as science driver




Analysis

- Performance seems to **saturate at ~ 40K** rollouts
- Still exploring only, a tiny fraction of the search space
 - Max learning depth of 472 (after 200k rollouts) is **only the first 472 decisions out of ~10,000**
- Limited storage capacity and downlink opportunities constrain how many images may be taken


Next Steps:


More analysis. Multiple satellites, Active fire mission.



*Tradeoff b/w
planner
process time
and optimality*

More rollouts =>
more
reinforcement
learning (more
process time) =>
better target
selection





- Near real time retrievals are challenging with limitations of auxiliary data which can be used.
- We shall explore GNSS-R retrievals from commercial sources such as Spire, available through the NASA Commercial Smallsat Data Acquisition (CSDA) program.
- We shall start work on incorporating the GNSS-R derived data with Active fire simulation models.
- Continue to keep focus on minimizing the responsive time of the system and evaluate the tradeoff with the performance.

Questions?



vinay.ravindra@nasa.gov



BACKUP SLIDES





Intelligent Tasking & Modeling






Planning time scales linearly with the # of rollouts

target = 10km dim location





Intelligent Tasking & Modeling



Max learning depth:
Deepest level in the search tree with reinforcement learning stats